

Studies in Linguistics and Philosophy 103

Tadeusz Ciecierski
Paweł Grabarczyk *Editors*

The Architecture of Context and Context-Sensitivity

Perspectives from Philosophy,
Linguistics and Logic

 Springer

Studies in Linguistics and Philosophy

Volume 103

Series Editors

Cleo Condoravdi, Stanford University, Stanford, CA, USA

Louise McNally, Pompeu Fabra University, Barcelona, Spain

Zoltan Szabo, Yale University, New Haven, CT, USA

Editorial Board Members

Johan van Bentham, University of Amsterdam, Amsterdam, The Netherlands

Gregory N. Carlson, University of Rochester, Rochester, NY, USA

David Dowty, The Ohio State University, Columbus, OH, USA

Gerald Gazdar, University of Sussex, Brighton, UK

Irene Heim, Massachusetts Institute of Technology, Cambridge, MA, USA

Ewan Klein, ICCS, Edinburgh, UK

Bill Ladusaw, Cowell College, University of California, Santa Cruz, CA, USA

Terrence Parsons, University of California, Irvine, CA, USA

Studies in Linguistics and Philosophy publishes monographs and edited volumes that focus on issues related to structure and meaning in natural language, as addressed in the semantics, philosophy of language, pragmatics and related disciplines, in particular the following areas:

- philosophical theories of meaning and truth, reference, description, entailment, presupposition, implicatures, context-dependence, and speech acts
- linguistic theories of semantic interpretation in relation to syntactic structure and prosody, of discourse structure, lexical semantics and semantic change
- psycholinguistic theories of semantic interpretation and issues of the processing and acquisition of natural language, and the relation of semantic interpretation to other cognitive faculties
- mathematical and logical properties of natural language and general aspects of computational linguistics
- philosophical questions raised by linguistics as a science.

This book series is associated with the journal Linguistics and Philosophy: <http://www.springer.com/journal/10988>

The series editors welcome proposals that fit the description above. For more information on how to submit your proposal, please contact the assistant editor, Anita Rachmat, E-mail: Anita.Rachmat@springer.com

More information about this series at <http://www.springer.com/series/6556>

Tadeusz Ciecierski • Paweł Grabarczyk
Editors

The Architecture of Context and Context-Sensitivity

Perspectives from Philosophy, Linguistics
and Logic

 Springer

Editors

Tadeusz Ciecierski 
Institute of Philosophy
University of Warsaw
Warsaw, Poland

Paweł Grabarczyk 
Institute of Philosophy, University Of Lodz,
Lodz, Poland
IT University of Copenhagen
København, Denmark

ISSN 0924-4662

ISSN 2215-034X (electronic)

Studies in Linguistics and Philosophy

ISBN 978-3-030-34484-9

ISBN 978-3-030-34485-6 (eBook)

<https://doi.org/10.1007/978-3-030-34485-6>

© Springer Nature Switzerland AG 2020

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors, and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, expressed or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

This Springer imprint is published by the registered company Springer Nature Switzerland AG.
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Acknowledgements

The work on this volume was funded by National Science Center, Poland, grant under award number 2018/29/B/HS1/01868. We thank all the authors for their high-quality contributions and patience and the anonymous reviewers for comments and suggestions. Let us add that the first versions of most papers collected in this volume have been presented at the 1st Context, Cognition and Communication Conference held at the University of Warsaw.

The Editors

Contents

An Introduction: The Architecture of Context and Context-Sensitivity ...	1
Tadeusz Ciecierski and Paweł Grabarczyk	
Utterance and Context	15
María de Ponte, Kepa Korta, and John Perry	
Indirectness and Intentions in Metasemantics	29
Michael Glanzberg	
Speaker Intentions and Objective Metasemantics	55
Jeffrey C. King	
Speakers, Hearers and Demonstrative Reference	81
Palle Leth	
How to Say When: A Reichenbachian Approach to the Answering Machine Puzzle	97
Agustin Vicente and Dan Zeman	
Distributed Utterances	113
Mark McCullagh	
Demonstratives in First-Order Logic	125
Geoff Georgi	
De Se as Variable Binding: On Context-Sensitivity in Utterance Reports	149
Joan Gimeno-Simó	
How Can “I” Refer to Me? Banishing Monsters at the Source	179
David Kashtan	

Compositionality in Truth-Conditional Pragmatics 205
Adrian Briciu

Occasion-Sensitivity and What Is Said 227
Claudia Picazo

Context and Communicative Success 245
Joey Pollock

Truth and Context 265
Gerald Vision

Subsentential Speech Acts: A Situated Contextualist Account 283
Joanna Odrowąż-Sypniewska

Some Constraints on Contextualism About Modals 295
Daniel Skibra

Index 317

Contributors

Adrian Briciu Department of Philosophy and Communication Studies, West University of Timisoara, Timisoara, Romania

Tadeusz Ciecierski Institute of Philosophy, University of Warsaw, Warsaw, Poland

María de Ponte Institute for Logic, Cognition, Language and Information (ILCLI), University of the Basque Country, Donostia, Spain

Geoff Georgi Department of Philosophy, West Virginia University, Morgantown, WV, USA

Joan Gimeno-Simó Department of Philosophy, University of València, València, Spain

Michael Glanzberg Northwestern University, Evanston, IL, USA

Paweł Grabarczyk Institute of Philosophy, University of Lodz, Lodz, Poland
IT University of Copenhagen, København, Denmark

David Kashtan Edelstein Center for the History and Philosophy of Science, Hebrew University of Jerusalem, Jerusalem, Israel

Jeffrey C. King Department of Philosophy, Rutgers University, New Brunswick, NJ, USA

Kepa Korta Institute for Logic, Cognition, Language and Information (ILCLI), University of the Basque Country, Donostia, Spain

Palle Leth Department of Philosophy, Stockholm University, Stockholm, Sweden

Mark McCullagh Department of Philosophy, University of Guelph, Guelph, ON, Canada

Joanna Odrowąż-Sypniewska Institute of Philosophy, University of Warsaw, Warsaw, Poland

John Perry Department of Philosophy, Stanford University, Stanford, CA, USA
Department of Philosophy, University of California, Riverside, Riverside, CA, USA

Claudia Picazo Department of Philosophy, University of Granada, Granada, Spain

Joey Pollock ConceptLab, University of Oslo, Oslo, Norway

Daniel Skibra Department of Philosophy, Northwestern University, Evanston, IL, USA

Agustin Vicente Ikerbasque Foundation for Science, Bilbao, Spain
Linguistics Department, University of the Basque Country (UPV/EHU), Leioa, Spain

Gerald Vision Temple University, Philadelphia, PA, USA

Dan Zeman Department of Analytic Philosophy, Slovak Academy of Sciences, Bratislava, Slovak Republic

An Introduction: The Architecture of Context and Context-Sensitivity



Tadeusz Ciecierski  and Paweł Grabarczyk 

Abstract In this paper we briefly present three perspectives from which the phenomenon of context-sensitivity in language can be studied: structural, functional and representational. Next, we describe the content of all papers included in the volume and the relations of the papers to the three perspectives in question.

It is a truism that context, content, and meaning interact. On the one hand, context determines what is communicated in the discourse; on the other hand, what is communicated in the discourse shapes the situation in which the discourse is embedded. Philosophers who agree about these basic facts disagree about nearly everything else: the scope of context-dependency, the nature of content, the character of input data of cognitive processes relevant for the interpretation of utterances, the role of beliefs and intentions as contextual factors, the validity of arguments in context-sensitive languages, etc. The growing interest in context-sensitivity and the development of sophisticated formal methods of representing discourse dynamics seems to stimulate diversity of opinion rather than consensus. There is no denying that part of this diversity comes from the sheer complexity of the studied phenomena. However, there is a lot of progress to be made toward unification and consensus. A closer look at the foundational philosophical issues regarding context-sensitivity might prove very fruitful in this respect.

Taking such a closer look is the aim of this volume. This will be done by considering at least the three different, yet related, perspectives on the phenomenon of context-dependence: *structural*, *functional*, and *representational*.

T. Ciecierski (✉)

Institute of Philosophy, University of Warsaw, Warsaw, Poland

e-mail: taci@uw.edu.pl

P. Grabarczyk

Institute of Philosophy, University of Lodz, Lodz, Poland

IT University of Copenhagen, København, Denmark

e-mail: pawg@itu.dk

© Springer Nature Switzerland AG 2020

T. Ciecierski, P. Grabarczyk (eds.), *The Architecture of Context and Context-Sensitivity*, Studies in Linguistics and Philosophy 103, https://doi.org/10.1007/978-3-030-34485-6_1

Inquiries into the *structural perspective* focus on the structure of contexts. Within this perspective, one might study the ontology of contexts, relations between various aspects of contexts, the nature of particular aspects of contexts, etc. Although some authors¹ have claimed that structural questions are constantly ignored in the literature, this is definitely not the case. Originally, it was Yehoshua Bar-Hillel (1954) who suggested that we think of utterances in terms of ordered pairs of expressions and contexts. Authors like Richard Montague, Dana Scott, and David Lewis proposed we think of the latter element in terms of “packages of miscellaneous factors relevant for determining extensions” (Lewis (1970): 24), that is, *contextual parameters*. Soon Hans Kamp, Arthur Prior, David Kaplan, and Robert Stalnaker showed that *contextual parameters* thusly conceived merge two functionally different roles: the content- (intension-) determining role and the extension-determining role. Thus, the opposition between *the context of utterance* (factors relevant for determining content/intension) and *the circumstances of evaluation* (factors relevant for determining extension) was brought to philosophical attention. Today, this distinction remains the point of departure for discussions concerning the phenomenon of context-sensitivity.

One of the first problems pointed out in that discussion (see: Cresswell 1973, 1979) was the *indefiniteness problem*. To make the concept of context-sensitivity broad enough, one needs to deal with an extremely rich list of contextual parameters. It seems that all properties of utterances can potentially matter for the utterance interpretation. The original suggestion due to Cresswell was to think of contexts not as a list of definite contextual parameters but rather as sets of arbitrary properties of utterances. This suggestion might be interpreted in several ways. One is as follows. The linguistic meaning of an expression-in-context (not to be confused with its content) determines which of the features of the utterance are relevant for its interpretation. Next, the meaning again determines how the values of the parameters are related to the semantic values of the lexical items that figure out in the utterance. This interpretation fits the behaviour of narrowly conceived indexicals, but it is rather hard to defend in other cases where it is difficult to make sense of the idea of linguistic meaning as determinant of the relevant contextual features. Take, for instance:

[1] The coffee can be found downstairs.

There is no question that the context somehow determines that, on a particular occasion, we are talking about coffee to drink during the conference break rather than, for instance, coffee beans to buy in the store downstairs (or *vice versa*). But nothing in the meaning of the sentence type [1] determines what features of respective utterances play the relevant content-determining roles. This might result

¹“Context is one of those linguistic terms which is constantly used in all kinds of contexts but never explained.” (Quasthoff 1994: 730).

in clearly distinguishing the semantic content and the utterance content. The first, the proper subject of semantics, must be linguistically mandated while the second, the proper subject of pragmatics or even specifically speech act theory, does not have to. This way of thinking might attract theorist who consider themselves to be *minimalists*.²

One alternative suggestion might be as follows. The linguistic meaning jointly with the background information of the participants of the conversational situation (information that, among other things, indicates what is relevant at a given stage of the conversational situation) determine which of the features of the utterance are relevant for its interpretation. Such analysis is friendly toward the idea of representing contexts as context sets (Stalnaker 1975, 1978, 1998, 2014),³ that is, classes of possibilities that are compatible with presuppositions of the participants of the conversational situation. Within this approach, one might try to distinguish between the types of processes (more or less linguistic meaning or background knowledge mandated) that are involved in the utterance interpretation process. From this point of view, one may suggest, the entire idea of the semantic content looks problematic as reducible to the idea of the utterance interpretation given certain assumptions about the conversational situation. This way of thinking might attract theorists who consider themselves to be *radical contextualists*.⁴

Seen from the *functional perspective*, contexts not only have an internal structure, but they primarily play various *roles* or, to put things differently, they can be *used in various ways*. One might study these uses as well as wonder about the demarcation line that separates context sensitivity and context insensitivity. Such investigations might result in functional typologies of contexts.⁵ One such typology of uses of context has been proposed by John Perry, who distinguished *presemantic*, *semantic*, *content-supplemental*, and *pragmatic* uses of context (Perry 2012: 47–59). Roughly speaking, the context is used presemantically if and only if it is used to determine which word or complex expression is actually used in cases where there is more than one candidate for being the word or the complex expression in question (like in the cases of lexical or syntactic ambiguities). A context is used semantically if and only if the meaning of a relevant expression (already fixed, if needed, in the process of

²Cf. Borg (2004b, 2012), Cappelen and Lepore (2005).

³Part of the motivation behind the idea of the context sets is dealing with the indefiniteness problem. As Stalnaker puts it: “When one looks beyond “I,” “here” and “now,” it is less clear how the facts about the situation in which the utterance takes place fix the relevant contextual parameters, and in the general case, the index theory, even in Kaplan’s version, says nothing about this. The pragmatic framework that I was promoting aimed to give a general representation of the features of a situation that determine the elements listed in an index.” (Stalnaker 1999: 6).

⁴Cf. Jaszczolt (2005) and Travis (2008).

⁵Such functional taxonomies should not be confused with non-functional ones. From the non-functional perspective one may, for instance, distinguish broad and narrow contexts as corresponding to more or less systematic context sensitivity (cf. Pelc 1971; Corazza 2004: 53–57; Bach 2012). One may also distinguish types of contexts depending on whether they meet certain descriptive constraints (cf. Stalnaker 2014).

presemantic use of the context) directs us toward particular aspects of circumstances of utterance enabling the assignment of semantic values to expressions (like in the cases of narrowly conceived indexicals). A context is used in a content-supplemental manner if and only if it determines the constituent of the proposition expressed (or, if you are not a friend of structured propositions, determines the proposition expressed) but without a mediating expression of morpheme present in the structure or the logical form of the utterance. This might happen in cases like:

[2] It is snowing.

In this case, there is no mention of the place, and the implicit reference to the place is not obligatory. However, the context somehow determines that the proposition expressed is that it is snowing at the particular place *p*.⁶ Last but not least, a context is used pragmatically if and only if it is used to determine the intentions of the speaker (like in the cases of implicatures or slips of the tongue).

Inquiries into the *representational perspective* embrace meta-theoretical questions such as: “What are methods of modelling context and context-dependence?” It has been suggested that at a given stage of the conversation, the speaker must presuppose that the information relevant for the interpretation of the utterance is available to other participants of the conversation. This fact, according to some authors, enables one to connect the parameter representation of contexts with the representation of the context as context sets because:

(...) the relevant contextual parameters must be available, and presupposed to be available, they will be incorporated into the speaker’s presuppositions, and so will be represented by the set of possible situations that constitutes the context set. (Stalnaker 1999: 6)

As this correspondence hypothesis illustrates within the representational perspective, one crucial question concerns the relation between various methods of representing contexts. Other possible considerations relevant for the representational perspective are: “What are the basic assumptions of formal semantic and pragmatic models of context-sensitivity?” (cf. Predelli 2005) “What are general desiderata for the empirically adequate theory of context dependence?” “What are the relations between language-qua-system-oriented and utterance-oriented approaches to context-sensitivity?” (the directival theory of meaning might be especially interesting in that respect, cf. Ajdukiewicz 1978; Ciecierski *forthcoming*; Grabarczyk (2019)).

The three perspectives are closely related: the functional perspective, for instance, cannot be separated from the structural one for any correct theory of the structure of contexts must take into account the roles contexts play or are expected to play (Perry’s taxonomy, for instance, might shed light on Cresswell’s indefiniteness problem). At the same time, the representational perspective is both shaped by functional and structural investigations and determines how one can describe the nature of contexts and methods that can be used in modelling the functions of

⁶For a discussion about the unarticulated constituents see also: Stanley (2000), Recanati (2002), Marti (2006), Hall (2008), Sennet (2011), Zouhar (2011), Valle (2018).

contexts. Hence it is important not only to address the representational, structural, or functional problems separately but also to study their mutual connections. The volume takes up all the issues described above while following this guideline.

In the first paper of the volume, Maria de Ponte, Kepa Korta, and John Perry compare two approaches to context-sensitive utterances: David Kaplan's theory (Kaplan 1979, 1989a, b), and the reflexive-referential approach (cf. Korta and Perry 2011; Perry 2012). They present consideration in favour of the latter approach. Roughly speaking, the difference between the two approaches might be seen as a difference in the way in which the notion of an utterance is construed. In Kaplan's theory, utterances (or rather *occurrences* as Kaplan prefers) are constructions consisting of expressions and contexts (just like in the case of Bar-Hillel's original idea from the 1950s) while on the reflexive-referential approach, utterances are basic notions that have particular roles assigned (contexts might be conceived here as constructed from such roles).⁷ As such, they remind Cresswell's representation of contexts as sets of relevant properties of utterances. On the reflexive-referential theory, one can distinguish at least two kinds of truth conditions of utterances (the two kinds of truth conditions, as the authors stress, are mediated by a large number of intermediate hybrid truth conditions). Referential truth conditions of utterances determine what else should be the case given the facts about the referential aspect of the utterance to make the utterance true. For instance, if Tadeusz utters "I am hungry now" on the 31st of December 2018 at 1 PM, the referential truth conditions shall be that Tadeusz has to have a certain property (namely: that of being hungry) on the 31st of December 2018 at 1 PM. The same referential truth conditions will be associated with the utterance "Tadeusz is hungry on the 31st of December 2018 at 1 PM." Reflexive truth conditions, on the other hand, are conditions taken from facts about the actual utterance and the contextual roles assigned to it. The reflexive truth conditions for "I am hungry now" are that the speaker of the utterance "I am hungry now" is hungry at the time of this particular utterance of "I am hungry now." Passing from contextual roles to occupants of contextual roles means passing from reflexive truth conditions to referential truth conditions. The introduction of utterances and reflexive and hybrid truth conditions, the authors suggest, enables one to distinguish various manners in which a given piece of information can be the subject of actions (like that of asserting) and cognitive states (like that of grasping). The authors show how this idea applies to various problematic cases, and they also consider its fruitfulness to be the main merit of the reflexive referential approach as nothing similar seems to be available in the case of Kaplan's account.

In "Indirectness and Intentions in Metasemantics," Michael Glanzberg presents⁸ the distinction between *direct* and *indirect* metasemantics. Roughly speaking, the semantic value of an expression is fixed in a direct way (so the expression and the

⁷The authors introduce in the paper a more general notion of episode. Utterances are special kinds of episodes. Others are, for instance, signals or thoughts. The idea of reflexive truth conditions extends to episodes thusly conceived.

⁸Originally introduced in Glanzberg's (2007).

corresponding contextual parameter have direct metasemantics) if it is determined by a single contextual factor. The semantic value of an expression is fixed in an indirect way (so the expression and the corresponding contextual parameter have indirect metasemantics) if it is determined by multiple contextual factors. The example that illustrates the case of indirect metasemantics is the case of adjectives like “tall” or “smart.” It is broadly assumed that the contextual parameter that determines the semantic values of adjectives is the standard relative to which the correctness of statements containing them is judged. Glanzberg argues that several factors fix the value of the standard parameter including at least comparison classes, presuppositions, overt discourse, and lexical meanings. Moreover, such factors might combine differently in different contexts. Next, Glanzberg argues for the indirect character of speakers’ intentions as contextual factors relevant for the interpretation of adjectives, especially in the context of a recent criticism of his view by Jeffrey King (2014). Firstly, he shows how King’s intentionalist account can be re-stated in terms of an indirect metasemantical account. Secondly, he argues for the claim that the direct intentionalist description of the semantics of adjectives is not the best way to go. One reason is that a variety of factors provide input to the cognition of standards. Another pertains to a broader picture of communication sketched by Glanzberg that, if accepted, supports indirect semantic account of standards. The next paper in the volume contains the reply of Jeffrey King. King presents briefly his idea of *supplementives*, that is, expressions that have semantic values in the context only if their invariant meanings are supplemented contextually. King proposed a single coordination semantics account for supplementives (cf. King 2013), which he summarises briefly in the paper. Then he carefully replies to Glanzberg’s arguments against a direct intentionalist account. Overall, the Glanzberg-King discussion is an excellent illustration of fundamental problems that arise within the perspective we have called *structural*.

In *Speakers, Hearers and Demonstrative Reference*, Palle Leth argues against the notion of objective demonstrative reference. The existence of such an objective reference is an assumption shared by most of the participants of the discussion about the contextual factors relevant for the semantics of demonstratives. Leth briefly describes the two main positions in the discussion: intentionalism, which claims that a speaker’s referential intention is the relevant contextual factor, and anti-intentionalism, which denies this and stresses the semantic importance of inter-subjective contextual factors that are accessible to hearers (both intentionalism and anti-intentionalism, as Leth stresses, have numerous versions and variants). Next, he criticises the participants of the debate for requesting a definite answer to the question about the referent of a demonstrative in some scenarios used in the debate (like Kaplan’s famous example of the Spiro Agnew picture (Kaplan 1978 or Reimer’s office keys (Reimer 1991)). He suggests, instead, that we should pay more attention to the hearer’s interpretative attitude toward demonstrative utterances. A careful inquiry into interpretative attitude reveals that the question regarding the objective reference of a demonstrative should be replaced with three other questions: the first regarding the speaker’s actual intention to refer, the second regarding the most reasonable assignment of reference, and the third regarding potential

(imagined) referential achievements of the speaker. Leth suggests that the issue of which question is relevant on a particular occasion depends on the interests of the hearer.

The next paper in the volume due to Agustin Vicente and Dan Zeman contains an exposition of an original approach to the answering machine puzzle (cf. Sidelle 1991; Cohen 2013) and other problematic uses of temporal indexicals. Roughly speaking, the puzzle concerns all uses of indexicals in general and temporal indexicals in particular where the relevant contextual parameter (in the case of temporal indexicals, the time of utterance) differs from the actual contextual parameter associated with the use of the sentence. The authors approach the problem using the ideas of Hans Reichenbach (1947) originally applied to grammatical tenses but extended by the authors to adverbs of time. In the Reichenbachian framework, three points in time are distinguished: the utterance time (UT), the event time (ET), and the reference time (RT). The combinations of the three enabled Reichenbach to describe the logical space of schemes of temporal reference (possible languages can partially or totally encode such schemes). Vicente and Zeman supplement Reichenbach's original framework with the idea that the entire frame of temporal reference can be moved (so it can take different points as the origins of the entire reference frame, that is, points that do not have to be identical with UT). This allows them to solve the answering machine puzzle and to accommodate other non-standard uses of temporal indexicals. The paper might attract the interest of tense logicians as it might be seen as an argument in favour of the Reichenbachian approach to the representation of tenses in language that is often considered now as inferior to its alternatives (Ohrstrom and Hasle 1995: 162–166).

Mark McCullagh investigates in his paper the idea of a *distributed utterance*, that is, “utterances that are more than usually spread out in time, location or other contextual features, in such a way that makes a difference to their proper interpretation.” For instance, one might utter a sentence containing two “heres” while walking and use the first here to refer to one location and the other to refer to the another one. On Kaplan's traditional account of indexicals, such intuitive interpretations are precluded as the single context is associated with entire sentences. Among other things, this requires that a single contextual parameter governs all uses of the same indexical type in the sentence (let us note that the problem might also concern uses of *distinct* but related indexicals in a single sentence or utterance like “today” and “yesterday”). McCullagh's proposal is to modify the semantics for the logic of demonstratives by introducing the distinction between simple and complex contexts, the former used for the evaluation of simple sentences only and the latter represented as n-tuples of simpler contexts, for the evaluation of complex sentences. One consequence of the modified approach is that in some contexts, formulas are not interpretable. McCullagh closes the paper with a discussion of possible applications of the modified semantics. They comprise the case of temporally and spatially distributed utterances, agentially distributed utterances, scare-quoting used in taking words in the other speaker's context, and solving Kit Fine's “antimony of the variable.”

In *Demonstratives in First Order Logic*, Geoff Georgi addresses the problem of multiple occurrences of demonstratives in a single sentence or, as he prefers, the problem of *referential promiscuity*, that is, the issue of providing a formal theory of expressions that might occur several times in a sentence while referring to distinct objects.⁹ Georgi argues that Kaplan's logic of demonstratives (LD) is an inadequate tool for intuitively correct analysis of validity in languages that contain referentially promiscuous expressions. There are at least two reasons for this. Firstly, certain sentences (as formalised in LD) are valid despite the fact that some of the utterances of such sentences might be false. Secondly, the idea of validity for indexical languages should be relativised to *particular contexts* rather than to *all contexts*. Georgi sketches a system of the first order logic of demonstratives that provides a tool for catching the intuitively correct idea of validity. In order to achieve this goal, he introduces the idea of context's coordination scheme and the relation of being a coordinated variant (applicable to pairs of contexts). Together with building into the semantics the occurrence-tracking parameters, the framework enables the revitalisation of the notion of validity to appropriately related contexts and classes of formulas. Contexts, on Georgi's account, have an interesting dual theoretical role. They are subject to generalisation when properties of validity and logical consequence are described and defined, but they also fix the coordination relations that remains fixed in the required generalisation. After presenting the formalisation of the modified logic of demonstratives, Georgi extensively discusses the relation of his account (and the idea of logic of demonstratives in general) to the assumption of formality of logic.

In the next paper, Joan Gimeno-Simó also addresses the problem of adequacy of Kaplan's logic of demonstratives. This time, however, the issue is not distributed utterances or multiple occurrences of indexicals but indirect speech reports (he is interested in *de se* reports, but he argues independently that the problem of *de se* is no harder than the traditional Frege puzzle). After presenting the basic features of Kaplan's original system, Gimeno-Simó argues that traditional semantics has resources for reconciling three claims. The first is an assumption about the syntax of attitude verbs, the second is a version of a principle called by Evans the intuitive criterion of difference for senses (Evans 1982:18–19), and the third is the claim that the only circumstances on which the truth value of the proposition depends are worlds and times (like in Kaplan's original system). His approach postulates a difference in the logical form of *de se* and *de re* belief reports. In the former case, the variable is obligatorily bound by the subject term (e.g., a proper name treated as a quantifier) while in the latter, it might be a free variable (a subject term treated as a regular singular term). The only (possibly serious) departure from Kaplan is that *de se* attitude reports contain monster operators. Gimeno-Simó applies his theory to

⁹This problem might be seen as a specific version of the problem of distributed utterances, so there is an interesting connection between Georgi's and McCullagh's papers.

the analysis of numerous examples, and he carefully discusses possible objections it must face.

The concept of a *monster operator* is the subject of the next paper in the volume. David Kashtan in *How can “I” refer to me? Banishing monsters at the source* argues for several claims. Firstly, he gives an argument for the general incompatibility of monsters and the idea of the (adequate) *formal* semantics for indexicals. Secondly, he proposes a revision of Kaplan’s system to comprise the incompatibility in question. Achieving both goals requires a clear distinction between *descriptive* and *pure* semantic endeavours. Most authors interpret Kaplan’s prohibition of monsters as a descriptive claim about natural languages (as such, it might be a subject of empirical falsification as suggested in Schlenker 2011). Kashtan’s analysis goes in the opposite direction; he proposes to treat it as a pure and foundational claim about formal systems, at least systems that meet additional criteria, that is, that are intended to capture the phenomenon of direct reference. He argues that monsters are definable in Kaplan’s semantics, making it inadequate as a formal system capturing direct reference. Finally, he sketches an interesting formal alternative based on the notion of content-generating directive. The alternative excludes monsters “at the source.”

The next paper in the volume due to Adrian Briciu contains a critical analysis of the concept of weak compositionality as used in various projects developed within the (moderate) contextualists camp where, in recent years, the idea of truth conditional pragmatics has gained some popularity. Roughly speaking, the difference between strong and weak compositionality consists in the fact that, in the case of the latter, context-sensitivity of complex expressions need not be traceable to the context-sensitivity of some of its constituents. Such a possibility is usually explained in terms of context-specific top-down pragmatic processes that change the literal content into the context-specific content. Briciu argues that weak compositionality is not capable of fulfilling the roles it is expected to fulfil, which might be seen as an argument undermining the theoretical projects that make use of weak compositionality. In particular, he stresses the highly problematic character of the notion of weakly compositional meaning rule. Firstly, the alleged existence of such rules conflicts with the idea of semantic indetermination that motivates contextualism. Secondly, given the attitude and agent-relative character of such rules, they would make the calculation of the content a matter of abductive reasoning (cf. Borg 2004a).¹⁰

The concept of semantic underdeterminacy is also the subject of the paper by Claudia Picazo Jaque. She asks the very important question of whether the concept might be extended to non-linguistic representational devices such as mental representations or concepts. In the contextualistic camp, two conflicting answers to this question are given (cf. Borg 2012: 518–521); proper contextualism claims that

¹⁰Briciu’s paper also highlights the semantic relevance of the key assumption of Cresswell’s indefiniteness problem, that is, the claim that all context sensitivity might be handled by means of a fixed and finite list of contextual parameters.

the answer is negative while occasionalism claims it is positive.¹¹ Claudia Picazo Jaque describes both views and presents three challenges to proper contextualism. Firstly, she argues that contextualism is committed to the gap between what natural languages can encode and what they can refer to, that contextualism has to claim that we can refer to definite properties while having no linguistic items that encode them. Secondly, she argues that the gap extends to mental representations (concepts) and that theories of the ones like Carston and Searle are not helpful in bridging the gap. Thirdly, she argues that contextualism is committed to the incorrect claim that properties correspond to total functions from objects to truth values. She claims that contextualism has no resources to meet the three challenges, and she closes the paper with sketching the occasionalistic alternative to contextualism. She stresses that the latter view takes the linguistic meaning to be merely preparatory and non-identifiable with application conditions and that it treats words as having application conditions only in the context of broader activities. Both Claudia Picazo Jaque's and Briciu's papers jointly constitute a strong case against moderate versions of contextualism.

In *Context and communicative success*, Joey Pollock develops an original approach to the idea of communicative success. Existing approaches to the problem of communicative success are, according to Pollock, committed to the assumption that standards for communicative success remain constant across different speech exchanges. Pollock contrasts such invariantism with a novel contextualistic idea of communicative success that assumes that standards of communicative success might differ across communicative situations. To use one example from the paper: the utterance "He is eating shellfish" might count as successful or unsuccessful in the situation where the addressee is unable to identify the referent of "he" depending on whether the speaker's intention was to have a certain effect on the addressee's actions (like causing the addressee to avoid eating shellfish) or just to say something about a particular person. Pollock provides several other examples of standard context sensitivity and identifies the two following contextual factors that play a role in setting the actual standard: the perlocutionary intentions of the speaker and the interlocutors' background beliefs. Pollock's account provides a good illustration of cases of Perry's pragmatic use of context that goes beyond implicature or slips of the tongue.

The next paper in the volume concerns the relation between truth and context with a special interest in parafictional statements, that is, statements about fictional states of affairs made outside fiction by actual persons. Firstly, Gerald Vision introduces the term *T-contextualism* to refer to theories that assume that context is a determinant of the truth-value of at least some statements and that a difference

¹¹One must be cautious here about the terminology. Some prefer to treat occasionalism as the proper version of contextualism. Also, the tag "truth-conditional" pragmatics is used slightly differently in the papers of Briciu and Picazo Jaque. The former author uses it as applicable only to views that assume weak compositionality while the latter uses it in the sense of the theory that claims that semantics and syntax are insufficient for determining the content of an utterance.

of context can make a difference to a statement's truth-value. In other words, T-contextualism, according to Vision, is interested in the role context plays in determining potential truth-makers of utterances. Such characteristics comprise both the theories in which the influence of context on truth value might be direct (like non-indexical contextualism) and those in which it is always indirect, that is, mediated by the proposition expressed (indexical contextualism). Secondly, he critically discusses a recent view of fiction by Timothy Crane (2014). Thirdly, he defends in the paper the conclusion that "if there are propositions whose subjects are fictional, context also makes a difference for their evaluation." He also defends the claim that parafictional statements are truth-evaluable relative to some contexts (but not relative to every context).

In *Subsentential Speech Acts: A Situated Contextualist Account*, Joanna Odrowąż-Sypniewska presents an approach to subsentential speech acts in the spirit of moderate relativism. Subsentsential speech acts are uses of expressions where speakers "produce projections of lexical items – which, seen semantically, are not of semantic type <t> and contain no force indicator – and yet they thereby make assertions, ask questions, etc. Since there is an assertion of something of semantic type <t>, what the speaker means in these cases extends beyond what her words mean" (Stainton 2004: 270). To account for subsentential speech acts, Odrowąż-Sypniewska uses the framework of moderate relativism developed by Recanati (2007). Roughly speaking, the idea is that the situation in which the utterance takes place can supplement the subsentential speech act by making it a full-blooded illocutionary act. However, while Recanati suggests that the analogous supplementation operates entirely at the level of circumstances of evaluation, Odrowąż-Sypniewska proposes (following the suggestion of Perry 1994) that it operates both at the level of the context of utterance and at the level of the circumstances of evaluation. She proposes also to use the disquotational indirect reports test to distinguish the context and circumstance-related contribution. For instance, we might correctly report someone's utterance of "It is snowing" as "She said that it is snowing," so the unarticulated constituents are elements of the circumstances of evaluation. At the same time, we cannot report someone's subsentential utterance of "From Spain" without explicit mention of the object (letter, suit etc.) that is from Spain. This suggests that the unarticulated constituents appear at the level of the context of utterance. Odrowąż-Sypniewska closes the paper with a slightly surprising conclusion that the very phenomenon of subsentential speech acts as well as her approach do not support contextualism. She conceives her view as a theory of speaker's meaning and not as a theory of what is said. Given this, the theory is consistent both with contextualism and minimalism.

In the last paper of the volume, Daniel Skibra discusses the relation between contexts and uses of modals. He starts the paper by distinguishing two dimensions across which the interpretation of modals might differ. The first, the *flavour diversity*, is connected with the type of modality that is relevant for the interpretation of the modal expression (deontic, epistemic, bouletic, teleological, etc.). The second, the *perspective diversity*, operates within each type of flavour of modality and provides additional information regarding the interpretation of the modal (like

whose knowledge or goals is relevant, what norms are at stake, etc.). Skibra briefly presents the idea of a uniform semantics for modals proposed by Angelika Kratzer and argues that one of its consequences is that “context alone is responsible for determining the flavour of a given occurrence of a modal.” Nonetheless, Skibra argues, we might show cases where certain interpretations of modals are not available. For instance, no epistemic interpretation is available for “must” with a preajcent containing a bare eventive predicate while additional considerations exclude the explanation of the fact that appeals to the idea of selectional restrictions. Next, the author sketches the alternative approach to modals (inspired by the theory of Valentine Hacquard; cf. Hacquard 2010). One crucial feature of the alternative is that the grammatical environment of the event variable (the introduction of event variables in the logical form of modal sentences is an important feature of Hacquard’s approach) can determine part of the modal flavour. The result is a restricted contextualism about modals.

References

- Ajdkiewicz, K. (1978). Language and meaning (1934). In J. Giedymin (Ed.), *The scientific world-perspective and other essays, 1931–1963*. Dordrecht: Springer.
- Bach, K. (2012). Context dependence. In M. García-Carpintero & M. Kölbel (Eds.), *The continuum companion to the philosophy of language* (pp. 153–184). London: Continuum International.
- Bar-Hillel, Y. (1954). Indexical expressions. *Mind*, 63, 359–379.
- Borg, E. (2004a). Formal semantics and intentional states. *Analysis*, 64(3), 215–223.
- Borg, E. (2004b). *Minimal semantics*. Oxford: Calrendon Press.
- Borg, E. (2012). *Pursuing meaning*. Oxford: Oxford University Press.
- Cappelen, H., & Lepore, E. (2005). *Insensitive semantics*. Oxford: Blackwell.
- Ciecierski, T. (forthcoming). *Indexicality, meaning, use*.
- Cohen, J. (2013). Indexicality and the puzzle of the answering machine. *The Journal of Philosophy*, CX, 5–32.
- Corazza, E. (2004). *Reflecting the mind*. Oxford: Oxford University Press.
- Crane, T. (2014). *The objects of thought*. Oxford: Oxford University Press.
- Cresswell, M. (1973). *Logic and languages*. London: Methuen.
- Cresswell, M. (1979). The world is everything that is the case, [w:] *The possible and the actual. Readings in the metaphysics of modality*, Loux (red.), Cornell University Press, s. 129–145.
- Evans, G. (1982). *Varieties of reference*. New York: Oxford University Press.
- Glanzberg, M. (2007). Context, content, and relativism. *Philosophical Studies*, 136, 1–29.
- Grabarczyk, P. (2017). Directival theory of meaning resurrected. *Studia Semiotyczne*, XXXI(1), 23–44.
- Grabarczyk, P. (2019) *Directival theory of meaning: From syntax and pragmatics to narrow content*, Synthese Library 409, Springer International Publishing.
- Hacquard, V. (2010). On the event relativity of modal auxiliaries. *Natural Language Semantics*, 18, 79–114.
- Hall, A. (2008). Free enrichment or hidden indexicals? *Mind & Language*, 23(4), 426–456.
- Jaszczolt, K. M. (2005). *Default semantics*. Oxford: Oxford University Press.
- Kaplan, D. (1978). Dthat. In P. Cole (Ed.), *Syntax and semantics* (pp. 221–243). New York: Academic.
- Kaplan, D. (1979). On the logic of demonstratives. *Journal of Philosophical Logic*, 8(1), 81–98.

- Kaplan, D. (1989a). Demonstratives. In J. Almong, H. Wettstein, & J. Perry (Eds.), *Themes from Kaplan* (pp. 481–563). Oxford: Oxford University Press.
- Kaplan, D. (1989b). Afterthoughts. In J. Almong, H. Wettstein, & J. Perry (Eds.), *Themes from Kaplan* (pp. 565–614). Oxford: Oxford University Press.
- King, J. C. (2013). Supplementives, the coordination account and conflicting intentions. *Philosophical Perspectives*, 27, 288–311.
- King, J. C. (2014). The metasemantics of contextual sensitivity. In A. Burgess & B. Sherman (Eds.), *Metasemantics: New essays on the foundations of meaning* (pp. 97–118). Oxford: Oxford University Press.
- Korta, K., & Perry, J. (2011). *Critical pragmatics: An inquiry into reference and communication*. Cambridge: Cambridge University Press.
- Lewis, D. (1970). General semantics. *Synthese*, 22, 18–67.
- Marti, L. (2006). Unarticulated constituents revisited. *Linguistics and Philosophy*, 29(2), 135–166.
- Ohrstrom, P., & Hasle, P. (1995). *Temporal logic: From ancient ideas to artificial intelligence*. Dordrecht: Kluwer Academic Publishers.
- Pelc, J. (1971). *Studies in functional logical semiotics of natural language*. Berlin: Mouton.
- Perry, J. (1994). Davidson's sentences and Wittgenstein's builders. *Proceedings and Addresses of the American Philosophical Association*, 68(2), 23–37.
- Perry, J. (2012). *Reference and reflexivity* (2nd ed.). Stanford: Center for the Study of Language and Information.
- Preddelli, S. (2005). *Contexts*. Oxford: Oxford University Press.
- Quasthoff, U. M. (1994). Context. In R. E. Asher & J. M. Y. Simpson (Eds.), *The encyclopaedia of language and linguistics* (Vol. 2, pp. 730–737). New York: Pergamon Press.
- Recanati, F. (2002). Unarticulated constituents. *Linguistics and Philosophy*, 25(3), 299–345.
- Recanati, F. (2007). *Perspectival thought: A plea for (moderate) relativism*. Oxford: Oxford University Press.
- Reichenbach, H. (1947). *Elements of symbolic logic*. London: Dave Publications.
- Reimer, M. (1991). Demonstratives, demonstrations and demonstrata. *Philosophical Studies*, 63(2), 187–202.
- Schlenker, P. (2011). Indexicality and de se reports. In von Heusinger, Maienborn, & Portner (Eds.), *An international handbook of natural language meaning*. Berlin/New York: De Gruyter Mouton.
- Sennet, A. (2011). Unarticulated constituents and propositional structure. *Mind & Language*, 26(4), 412–435.
- Sidelle, A. (1991). The answering machine paradox. *Canadian Journal of Philosophy*, 21(4), 525–539.
- Stainton, R. J. (2004). The pragmatics of non-sentences. [w:] Horn, Ward (red.) *The Handbook of Pragmatics*, s. 266–287.
- Stalnaker, R. C. (1975). Indicative conditionals. *Philosophia*, 5, 269–286.
- Stalnaker, R. C. (1978). *Assertion, syntax and semantics 9*. New York: Academic.
- Stalnaker, R. C. (1998). On the representation of context. *Journal of Logic, Language and Information*, 7, 3–19.
- Stalnaker, R. C. (1999). *Context and content*. Oxford University Press.
- Stalnaker, R. C. (2014). *Context*. Oxford: Oxford University Press.
- Stanley, J. (2000). Context and logical form. *Linguistics and Philosophy*, 23, 391–434.
- Travis, C. (2008). *Occasion-sensitivity*. Oxford: Oxford University Press.
- Valle, R. (2018). Weather predicates, unarticulation and utterances. *Manuscrito*, 41(2), 1–28.
- Zouhar, M. (2011). What is wrong with unarticulated constituents? *Human Affairs*, 21(3), 239–248.

Utterance and Context



María de Ponte, Kepa Korta, and John Perry

Abstract In this paper we explain two approaches to context-sensitive utterances, focussing on temporal indexicals and tense. The first approach is David Kaplan's account in "On the Logic of Demonstratives" (1979) and "Demonstratives" (1989). The second is the reflexive-referential approach used by Korta and Perry in *Critical Pragmatics* (2011). We argue for the second approach, using a famous example of Arthur Prior's.

Keywords Context · Episodes · Indexical · Reflexive-referential · Tense · Utterances

1 Introduction

If JP¹ says, "I live in California," he says something true about JP. If MDP says, "I live in California," she says something false about MDP. Utterances of sentences with indexicals such as "I" express different propositions, depending on *context*, in this case the speaker. Similarly with tense. An utterance by KK of "I am tired," early Sunday morning as he gets out of bed, may be false. But 90 min later, returning from running his daily 5 K to get ready for mass, his utterance of the same words will be true.

¹We refer to ourselves in this paper as MDP, KK, and JP.

M. de Ponte (✉) · K. Korta
Institute for Logic, Cognition, Language and Information (ILCLI), University of the Basque Country, Donostia, Spain
e-mail: maria.deponte@ehu.eus; kepa.korta@ehu.eus

J. Perry
Department of Philosophy, Stanford University, Stanford, CA, USA
Department of Philosophy, University of California, Riverside, Riverside, CA, USA

In this essay we consider what seems to be an issue of detail. On Kaplan's approach, contexts are sets, quadruples of a speaker, time, location and world. Utterances do not appear in the theory, but are modeled by pairs of expressions and contexts. An expression has a *character* (meaning); an expression-in-context has a *content* (proposition or component thereof.)

On the account we favor, the "reflexive-referential" theory, utterances appear in the theory; they are what the theory is about. Speaker-of, time-of, and location-of are *roles*, that is, functions from an utterance to the object that stands in the appropriate relation. The term "context" is used with its usual meanings for the truth-conditionally relevant circumstances of an utterance, but contexts as such are not entities within the theory.

We argue that this difference is more significant than it might seem, and that our approach has advantages for understanding the relation between the content of utterances (and other contentful episodes), their causal roles and their cognitive significance. On the other hand, we do not deny Kaplan's point, that modelling utterances as pairs of expressions and context has advantages in developing a logic of indexicals and demonstratives.

2 Prior's Root Canal

Here is a famous quote from Arthur Prior

[H]alf the time I personally have forgotten what the date is and have to look it up or ask somebody when I need it for writing cheques, etc.; yet even in this perpetual dateless haze one somehow communicates, one makes oneself understood, and with time references too. One says, e.g. "Thank goodness that's over!" . . . says something which it is impossible that any use of a tenseless copula with a date should convey. It certainly doesn't mean the same as, e.g. "Thank goodness the date of the conclusion of that thing is Friday, June 15, 1954," even if it be said then. (Nor, for that matter, does it mean "Thank goodness the conclusion of that thing is contemporaneous with this utterance." Why would anyone thank goodness for that?). (Prior 1959: 17)

Consider two sentences Prior might have used to give thanks on June 15, 1954 at 7 p.m.

Thanks goodness that . . .

(1) the root canal is over [now].

(2) the root canal is over as of Friday evening, June 15, 1954.²

(1) seems a normal thing for a person who has just undergone a root canal to say. (2) does not, although with a bit of effort one can construct an example where it makes a bit of sense — perhaps one's dental insurance lapses on June 16. The cognitive and emotional significance of (1) and (2) — the doxastic and emotional states that

²Actually, June 15, 1954 was Tuesday, and not Friday, as Prior indicates. We will, however, keep Prior's example as it is.

would typically cause such utterances and of which such utterances would be signs — are different.

MDP and KK argue that if one supposes, as Prior does, that the difference in cognitive significance must be due to difference in the propositions expressed by (1) and (2), then metaphysical consequences loom. And Prior draws such conclusions. If we confine ourselves to objective facts about the temporal events — the “B-series” in McTaggart’s (1908) terminology — then we can give the truth-conditions of (2). But to get at the different truth-conditions of (1) we need “A-properties,” such properties as being past, present and future. Objective facts about insurance coverage dates might make one happy that one’s root canal did/is/will take place before June 15, 1954. But it is the diminution of fear and apprehension of *future* pain that makes one happy that ones root canal has *already* taken place.

MDP and KK point out that the metaphysical conclusion depends on taking the difference in cognitive significance to be a difference in proposition expressed.³ On either Kaplan’s approach or on our approach, there is a more plausible alternative; one can believe the same proposition in more than one way, and the different ways of believing account for differences in cognitive significance.

2.1 Kaplan’s Approach

In Kaplan’s theory, a context is a set, a quadruple of a speaker, a time, a location, and a world. A proper context is one in which the speaker is in the location at the time in the world. The meaning, or *character* of an expression is a function from contexts to *contents*. The content of a sentence is a proposition, that of other expressions is their contribution to the proposition expressed by sentences in which they occur, basically an object, property, or relation. Thus (1) and (2) have the same content in the circumstances we are imagining, that Prior’s root canal occurs prior to Friday evening, June 15, 1954.

In a series of essays, JP defended Kaplan’s theory, and argued that Kaplan’s concept of character was not only a contribution to understanding how indexicals and demonstratives work, but also to understanding intentionality in general.⁴ Basically, JP claimed that the causal roles of perceptual states, states of belief and

³See de Ponte and Korta (2017).

⁴See the early essays collected in JP’s *The Problem of the Essential Indexical and Other Essays* (2000), especially Perry (1977, 1979). The move to the reflexive-referential begins in the essay “Cognitive Significance and New Theories of Reference” (1988), with the concept of “the proposition created.” This essay was a reply to Howard Wettstein’s important essay “Has Semantics Rested on a Mistake?” (1986). In his essay, Wettstein introduced examples in which the expression, character, and context are basically the same, but the cognitive significance is nevertheless different. The reflexive-referential theory is developed in *Reference and Reflexivity* (2012 [2001]) and in Korta and Perry, *Critical Pragmatics* (2011). The notion of episode was first

desire, and intentions and volitions can only be understood in terms of character. JP's term for this generalized concept of character was "role". The causal role of a visual state, for example, is to carry information about the objects in view of the perceiver at the time and place of perception. The same state can carry information about different objects at different locations at different times, for different perceivers or the same perceiver at different times and places. We cannot understand this simply in terms of propositions that encapsulate the information captured, but must also bring in roles, functions from the circumstances the perceiver is in to *what* is perceived.

2.2 *The Reflexive-Referential Approach*

Our position differs somewhat from Kaplan's view, and the view that JP defended. Kaplan discusses utterances to motivate his theory, but he does not bring them into his theory as such; they are replaced by, or perhaps modeled as, pairs of expressions and contexts: "expressions-in-context". Kaplan's main interest was developing a logic of demonstratives and indexicals. For this purpose he regarded utterances as an unnecessary complication. For one thing, a pair of context and expression can have a content, even if the speaker of the context does not utter the expression at the time and place of the context in the world under consideration. More importantly, utterances take time; the validity of an argument with 100 steps might depend on the context being the same for all of them, but we can't talk or write that fast. So, for logical purposes, utterances can get in the way.

From the point of view of understanding the relations between the contents of states and their causal roles, however, it is very helpful to have *episodes*, — paradigmatically, utterances, but also perceptions, thoughts, and actions — in our theory, as well as more extended "episodes" such as beliefs, desires and intentions. It is such episodes that *have* contents, have causes, and have effects. So, in the theory of intentionality, episodes and, in particular, utterances are too important to ignore, in spite of the complications they pose for logic.

For this purpose, we consider the elements of Kaplan's contexts to simply be properties of utterances, which objects fill the roles of speaker-of, time-of, and location-of. The fact that utterances have speakers and occur in locations at times clearly inspired Kaplan's concepts of context and character. We promote these inspiring episodes to first-class status.

The chief advantage of our view is simply that that it accounts for — and makes use of — the fact that utterances have many other properties in addition to having

introduced in Crimmins & Perry, "The Prince and the Phone Booth: Reporting Puzzling Beliefs" (1989), in their account of beliefs as concrete cognitive particulars.

speakers, locations, and times, that can be relevant to understanding their cognitive significance.⁵

One way to handle these, while sticking with Kaplan's approach, is to add more members to the context set, or to introduce additional sets. The latter is more or less the approach of Jon Barwise and JP in *Situations and Attitudes* (1983). On the "relational theory of meaning" advocated there, the meaning of a sentence is taken to be a relation among various situations connected to an utterance, although the utterance itself is, as in Kaplan's theory, only modelled and not introduced directly into the theory. The basic relation is between the utterance situation, which determines the speaker, location and time, and, in lieu of propositions, described situations. But various other situations, are added to the range of the relation, for dealing with names, descriptions, ambiguity and other phenomena.

On the reflexive-referential account, however, the treatment of such factors is simpler and more straightforward. They are all properties of the utterance, which can be recognized as necessary to handle various phenomena.

On the reflexive-referential theory, truth and falsity are regarded as properties of episodes. An utterance has truth-conditions, and is true if it satisfies them. Take a simple example, JP's utterance *u* of "I am sitting". For *u* to be true, there must be a speaker of *u* and a time of *u*, and the speaker must be sitting at the time. These are the *reflexive* truth-conditions. Then, *given* that JP is the speaker and noon August 28th is the time of *u*, JP must be sitting at that time for it to be true. That is, by identifying the occupants of the roles in the reflexive truth-conditions, we obtain the *referential* truth-conditions. The latter are not, in themselves, conditions on the utterance. JP could be sitting at that time without uttering anything. But if we conceive of the referential truth-conditions as giving *what else* has to be the case for the utterance to be true, *given* the referential facts, the referential truth-conditions are conditions on the utterance.

In the referential-reflexive theory, we distinguish being true and being factual. Truth is a property of utterances and other episodes. Being a fact is a property of a state of affairs, or circumstances, or whatever else one takes to serve as *possibilities*.

That JP is sitting at noon is a fact, because he is. No utterance or episode is required. It would be a fact that Venus is the second planet from the sun, even if no one ever said so, even if there were no language, or even no life on earth, and no utterances. But for an utterance of "Venus is the second planet from the sun" to be true, life, language, and speakers are all required.

Thus the truth-conditions of each utterance are determined by the expressions used and the occupants of relevant roles. Utterances with different expressions, and different occupants of the relevant roles, will have different truth-conditions. Truth, the property that all true utterances have in common, is the property of meeting the truth-conditions that an utterance provides for itself.

Usually the phrase "truth-conditions" is used in contemporary philosophy for what we call "referential truth-conditions," what *else* has to be the case for the

⁵We consider the issue of possibilities and worlds in the last section.

utterance to be true, *given* the facts of reference. In this sense, quite different utterances can have the same truth-conditions, for example JP's utterance of "I like philosophy now" and KK's simultaneous utterance to JP of "You like philosophy now". But the reflexive truth-conditions will not be the same; the former requires the speaker of that very utterance to like philosophy to be true, the latter requires that the person the speaker of that utterance is addressing likes philosophy.

Let u be an utterance of "I love sailing now". Taking just the meaning into account, we can say:

u is true iff the speaker of u loves sailing at the time of u .

As noted, these are conditions on the utterance u , properties *it* must have to be a true; that is, conditions on the utterance *itself*, and hence *reflexive truth-conditions*.

Suppose MDP is the speaker of u , and July 2017 is the time. Given that, we can give the *referential* truth-conditions of u , that is, what else has to be the case for u , given the reference of 'I' and 'now'. Note that they put *no* conditions on u , but on MDP and July 2015. They are also the referential truth-conditions of KK's utterance at the same time to MDP, "You love sailing now!"

We argue that the reflexive-referential theory inherits a key insight of Kaplan's theory, and JP's earlier view, but the inclusion of utterances gives it two advantages. The inherited key insight is the distinction between different ways in which information can be discovered, believed and asserted.

The first advantage is that the cognitive significance of an utterance for different hearers can depend on the perceptual and causal relations the hearer has to the utterance, and which reference-determining referential facts they know, and how they think of them. To account for this, we need to bring in additional properties of the utterance, and in particular causal properties. These are not modelled by Kaplan's expressions-in-context.

The second advantage concern certain possibilities that are hard to find unless, again, we have utterances — or more generally cognitive and linguistic episodes — in our account. We elaborate on these two advantages below. We start with the first one, on Sects. 3 and 4, and we discuss the second one on Sects. 5 and 6.

One final point, before we leave Prior's example for a bit. On his list of things it would be strange to thank goodness for he includes

(3) The conclusion of the root canal is contemporaneous with this utterance.

The referential truth-conditions of (3) are that the conclusion of the root canal is contemporaneous with the utterance (3). This is about the same as the reflexive truth-conditions of (1).

(1) the root canal is over [now].

The reflexive truth-conditions of (1) are that the root canal is over at the time of utterance (1).

In general, an utterance that elevates the relevant reflexive truth-conditions of an original utterance to referential truth-conditions will not have the same cognitive significance as the original. Compare, "I need some salt," "KK needs some salt"

and “the speaker of this utterance needs some salt”, all said at dinner by KK. The relevant reflexive truth-conditions of the first are that the speaker of the utterance needs some salt. This is what KK intends to convey; that is, he wants others at the dinner table, who can reach the salt, to be able to identify the one who needs salt in a way that will lead them to pass the salt in the right direction. They hear the utterance; they can easily identify the speaker; they know where the speaker is in comparison to them, so they will know how to get the salt to the person that needs it. The second utterance does not have these virtues; it will help KK’s fellow diners to help the person who needs the salt only if they know who KK is; even if this can be assumed it sounds pretentious.⁶ The third utterance is better than the second. As long as his fellow diners know that KK is referring to his own utterance with “this,” they can figure out where to pass the salt. Otherwise they might wonder what utterance he is referring to — perhaps something he reads on the menu. By referring to his own utterance, rather than simply making it, KK makes the utterance part of the referential content, that is, part of what he is talking about. Even if they realize which utterance he demonstratively refers to, and manage to pass the salt in the right direction, this will strike his fellow diners as odd — perhaps a way of emphasizing KK’s obsession with utterances.

Similarly, in (3), Arthur Prior refers to his own utterance, instead of simply producing it, making the utterance itself part of the referential content, that is, part of what he is talking about. To understand him, the audience would need to know what utterance he is referring to; and even if they do, it will strike them as odd. Thus the cognitive significance of (3) — which talks about (1) or, in other words, includes (1) in its referential truth-conditions — is different from the cognitive significance of (1).

3 Varieties of Cognitive Significance

The first advantage mentioned above is that in the reflexive-referential theory there is a simple account of how utterances have have different cognitive significance for different people, depending on their relation to the utterance. To account for this, we need to bring in additional truth-conditions of the utterance; and additional properties. Utterances have reflexive and referential truth-conditions, as we noted. But many other levels of truth-conditions can be considered for different purposes. The reflexive content of MDP’s utterance *u* is simply that the speaker of *u* loves sailing at the time of *u*. Suppose MDP steps off the boat at a pier in San Francisco Bay after an afternoon of sailing. After accidentally falling in the cold water of San Francisco Bay during the sail, she had said to KK, “Sometimes I hate sailing.” But as she steps off the boat she utters *u*, “I love sailing now.” A stranger, who like Prior, is in a dateless haze, and is looking in the opposite direction, hears her. Initially,

⁶See Korta and Perry (2011), Chapter 7.

he has only an utterance-bound “cognitive fix” on the speaker of *u*: whoever is the speaker of the utterance he hears.⁷ Then he turns, and recognizes that the speaker of the utterance is the young lady he sees. Now he knows that she is the person who must now love sailing for the utterance he heard to be true. Also, the stranger originally has only an utterance-bound fix on the time of *u*: the time of the utterance he hears. Once he turns, and realizes that he is hearing the utterance at the same time it occurs — rather than hearing it over the radio — he realizes that the time of the utterance is also the time of his hearing of it; it is the time that, even in his dateless haze, he thinks of as “now.” So, as he gains more knowledge about the properties of the utterance, he moves from only grasping the reflexive truth-conditions to grasping the referential truth-conditions, which he can express with “She likes sailing now.”

Suppose you see a video of MDP, whom you recognize, uttering *u*, but have no idea when the video was taken. If you believe her, you would come to believe that at the time of the utterance she enjoyed sailing. Your understanding is utterance-bound, for you can identify the time at which she loved sailing only as the time at which *u* occurred. But since you recognize MDP, you have more than a reflexive understanding. The relevant truth-conditions are conditions both on the time of the utterance and on MDP. In *Critical Pragmatics*, KK and JP call such truth-conditions “utterance-bound,” and considered a large number of examples and issues in which utterance-bound truth-conditions (and other “hybrid” truth-conditions) are crucial to understanding semantic and, in particular, pragmatic issues.

4 Back to Prior

On both Kaplan’s approach and the reflexive-referential approach, Prior’s utterances (1) and (2) are importantly different in cognitive significance. On both approaches, the explanation for the difference need not rely on the proposition expressed, but in the way of apprehending or asserting that proposition. So either approach can avoid Prior’s metaphysical conclusions.

To see the advantages we claim for the reflexive-referential approach, let’s assume that Prior’s philosophical thoughts were inspired by a real root canal that he had on June 15, 1954, in early afternoon. Towards evening he went to a bar with friends, all of whom, like Prior, went around in dateless hazes. It hurt him to talk, but he wrote a note and showed it to his friends:

(4) I had a root canal earlier today.

The note is preserved in the Museum of Tense Logic in Auckland, together with a little explanation of how it came to be written. We examine it, more than half a century later. The cognitive significance of the note was much different for his friends at the bar than it will be for us.

⁷We own the phrase “cognitive fix” to Howard Wettstein.

Such a note is a token, the physical product of act of writing, which we take to be a species of uttering. When one takes a note to be meaningful, one thinks of it as the result of an intentional act of writing. The context for the expression-in-context, or the occupants of the relevant roles for an utterance based theory, is provided by the intentional act that produced the token.⁸

Thus the context, character, and content of writing the note don't change between 1954 and 2017. And it is the same utterance of the same expression and the same objects filling the utterance-relative roles and the same reflexive and referential contents, in 1954 and 2017.

Let's imagine that Prior's friends and the museum-goers both were in dateless but not totally clueless hazes. That is, they couldn't provide the date and time they became aware of the note, but they knew the year — 1954 and 2017 respectively.

Thus, when they inspected the note, both groups would understand its truth-conditions at an incremental but still utterance-bound level:

The display of this note was true, iff it occurred on a day when Prior had a root canal, later in the day than the root canal.

The difference is that the two groups, although both in dateless hazes, know different things about the utterance. Prior's friends, who witnessed the display, know that it occurred on the same day that they are sitting with Prior in the pub and can offer him a drink. The museum-goers, aware that Prior died in 1969, know that whatever the exact date of the utterance, or the exact date of their inspection, the first occurred many years before the second. And this difference explains why the first group offers Prior a drink, but the second group makes no such effort.

Kaplan's theory also has the resources to explain the difference. Kaplan's contexts contain not only speakers, locations, and times, but also possible worlds. So the context of Prior's utterance contains all the facts that could be possibly relevant, including all facts about his utterance, its effects in the bar, and much later in the museum. Even if we do not have utterances in our basic semantics, we can bring them in through the "back-door", by finding them in the possible world in the context.

To us, this seems a bit roundabout. The utterances, and other episodes the truth-conditions and cognitive significance of which we want to understand, are what our semantic and pragmatic theories are ultimately theories about. It seems natural to give them a central place in our theories. In addition, the inclusion of utterances in our theory brings a second advantage for the reflexive-referential theory over the Kaplan-Perry one.

⁸See Predelli (1998, 2011) and Perry (2003) for complications.

5 Wettstein's Challenge

In his important essay, “Has semantics rested on a mistake?” Howard Wettstein pointed out that whatever the virtues of what he called the “Kaplan-Perry” account has in explaining cognitive Significance of cases involving indexicals, it does not handle Frege’s *Begriffsschrift* (1879) problem, the origin of worries about cognitive significance, which involves proper names. “Hesperus is Hesperus” and “Hesperus is Phosphorus” clearly have different cognitive significance; one learns from the second that the names co-refer, but not from the first. On a directly referential account of proper names, which is more or less what Frege had in the *Begriffsschrift*, this is hard to account for. On Kaplan’s account, the character of a proper name is a constant function, from any context to the bearer of the name. The two sentences have the same content, a singular proposition to the effect that Venus is Venus. So, whatever the virtues of the Kaplan-Perry thesis for cases involving indexicals, it does not help with proper names.

In response to Wettstein’s essay, JP introduced the concept of the proposition created by an utterance, in contrast with the proposition expressed by an utterance, which was basically the distinction between reflexive and referential content.

On the reflexive-referential account, the two utterances have different reflexive truth-conditions. The first is true if and only if there is an object named by “Hesperus” which is self-identical. The second is true if and only if there is such an object, and there is also an object named by “Phosphorus”, and the objects are identical. So, even though the contents or referential truth-conditions of the two utterances are the same, they differ in cognitive significance, in virtue of their different reflexive truth-conditions.

6 More Hazes

As noted, certain possibilities are hard to find unless we have utterances —and other episodes —in our account. We turn now to the discussion of this second advantage.

One can be in a dateless haze without knowing it; that is, one can be quite certain about the date, but be wrong. In such a case, it is natural to think, in retrospect, that one *might have* been correct; even that one’s false belief was justified. This seems to provide another advantage for the reflexive-referential theory.

Suppose JP and Dan are planning to go to the Giants game on August 22, 2017. The day before, JP types out a reminder: “The Giants game is tomorrow. Don’t forget.” But he forgets to hit the “send” button. He notices that that the message has not been sent just before retiring, and hits the button. But he doesn’t notice that it is already after midnight.

Dan, a dateless-hazer, sees the reminder, “The Giants game is tomorrow,” when he wakes up on August 22, and sees that date on the email heading. He knows that the game is on August 22, and reasons, given JP’s notorious reliability, “Today

must be August 21.” He then immerses himself in linguistic esoterica until late in the evening, when JP calls and says, “You missed the game!”

Dan thought, on August 22, that it was August 21. He had good reason for this belief. The game was scheduled for August 22. JP said, in an email this morning, that the game was tomorrow. JP is pretty reliable. Therefore, today must be August 21.

Could Dan have been right? It seems not, because for his belief to be true, August 22 would have to be August 21, which is not possible. Surely had Dan said,

(5) August 22 is August 21

we could diagnose some kind of irrationality (or was making some subtle linguistic point). But if he just says, as he did:

(6) Today is August 21

on August 22, this doesn’t seem like the right diagnosis.

But there is a way Dan’s utterance could have been true: if it had occurred on August 21 rather than August 22. That is, specifically, if the role of time-of for the utterance of ‘today’ had been filled by some moment occurring on August 21, the utterance would have been true.

On the expression-in-context approach, this doesn’t seem like an option. Since the pair of expression and context is individuated by its members, we wouldn’t have the same pair if the time of the context was August 21. So we have another advantage for the reflexive-referential theory and granting utterances first-class status in the semantics of tense and indexicals.

As Richard Vallée has pointed out,⁹ there is an objection to our strategy. One theory of events is that they are *individuated* by the time at which they occur, *where* they occur, and *which* object and properties involved. If we accept this account of event individuation, the reflexive-referential account is no better off than the expressions-in-context account.

We reject this account of the individuation of events; it is a plausible account of the individuation of facts, but not of events. Being a fact is, as noted, a property of whatever one takes to serve as possibilities, be it circumstances, state of affairs or whatever. We will not here develop an account of events, which we regard as very basic elements of reality. But we think an adequate account must allow for counterfactuals of the form, “if the election had occurred two weeks earlier, Clinton would have won,” or, to follow with our example, “if Dan’s utterance had occurred on August 21, he wouldn’t had missed the Giant’s game.” It is a fact that he missed it and that he made his utterance on August 22. Nothing can change that. But the episodes involved, Dan’s belief and Dan’s utterance “Today is August 21” could have been true, had they occurred on August 21.

In his paper “Frege on Demonstratives” (1977), JP introduced the example of Heimson, who thought he was David Hume. Let’s suppose instead that Heimson

⁹Personal Communication.

thought he was Bob Dylan, which will make it easier to make a case for his rationality. Here is the background story. Heimson falls, hits his head, and has amnesia as a result. He doesn't know who he is. He carries no identification. He awakes in a hospital where no one has any idea who he is. Heimson decides to figure out who he is. Heimson's amnesia is of a rather peculiar sort; he retains "third-person" memories about lots of people, he simply doesn't remember which of them he is. He assumes he is one of the people about whom he knows a great deal. He notices that he knows all of Bob Dylan's songs by heart, the date of every concert where he performed, and loads of other things. He also knows a lot about a fellow named "Heimson," but not nearly as much as he knows about Dylan. He decides he is Dylan, and thinks, with some confidence, "I am Bob Dylan".

His thought cannot be true. Indeed, it seems necessarily false. But it might be rational. And we think it does get at a possibility, even if a rather remote one. The possibility is found at the reflexive level. Call his thought — the event of thinking, the episode, not its content — T. If Bob Dylan had the thought T, rather than Heimson, T would be true. This is the possibility that Heimson's sifting of the evidence available to him led him to think was the case.

David Lewis (1979), considering JP's original example, comes to the opposite conclusion. In "Frege on demonstratives," JP advocated a version of the Kaplan-Perry view, that he called the two-tiered view. In a nutshell, to deal with the attitudes we need to recognize two levels of content for beliefs. *What* is believed is a proposition, often a singular proposition. *How* it is believed corresponds to character or role.

Lewis had nice things to say about this account, but thought that the level corresponding to character or role could serve as what is believed, and the upper tier could be jettisoned. Lewis noted that a character, a function from contexts to propositions, could be regarded as a property: the property an agent has at a time iff the character, applied to that agent and time, yields a true proposition. So, when JP says, "I am sitting," he "self-ascribes" the property P_{sitting} :

P_{sitting} =
 the property x has iff the character of "I am sitting", with arguments x and t, yields a true proposition, that is,
 the property someone x has at time t, iff x is sitting at t.

On Lewis view, properties, rather than propositions, are the true "objects" of beliefs. A belief consists of an agent at a time self-ascribing a property.

Lewis's view, like the Kaplan-Perry approach, does not involve episodes. We have agents, times, the relation of self-ascription, and properties. Lewis (1979, footnote 16) regards singular propositions and "de re" beliefs as unnecessary intrusions into the theory of the attitudes based on pre-occupation with the analysis of our customs for reporting beliefs.

Lewis's view incorporates much of the traditional picture that belief is a relation between an agent at a time and a proposition. His innovation is to replace propositions with properties. But we think the idea of objects of beliefs in this sense is a mistake. A belief consists in an agent at a time being in a mental state, an episode. This episode has truth-conditions, which can be characterized

by propositions, many different propositions, depending on what is taken as given. But neither propositions, nor characters, nor characters construed as properties, are objects of belief in the sense that belief consists of a relation to them. Propositions, in our sense, are tools we use to characterize and keep track of the truth-conditions of the episode.

Lewis's account, like the Kaplan-Perry account, does not have episodes, utterances or beliefs, as elements. So, on Lewis's view, Heimson's belief, whether he thinks he is Hume or thinks he is Dylan, cannot be true in the strong sense that there is no possible way it could be true. On his view, the belief consists of Heimson, the relevant time, and the attitude of self-ascription to the property of being Hume/being Bob Dylan. Since there is no possibility that Heimson at the time has that property, there is no way the belief can be true. Heimson is irrational in all possible circumstances.

So, we count it as a defect of the Kaplan-Perry account, and of Lewis's account, that it is unfair to Heimson. To paraphrase Billy Joel, Heimson may be crazy, but often in philosophy we are looking for a lunatic, to uncover hidden corners in the realm of possibilities.¹⁰

7 Conclusion

The topic of this volume is contexts in general. We do not hold that theories that treat contexts as abstract entities, sets that encode relevant contextual information, are necessarily wrong-headed. For the purposes of formal theories that lend themselves axiomatization and computation, such theories have many virtues. But we think that in order to fully understand what is going on for many philosophical purposes, it must be kept in mind that possession and transfer of information are always a matter of complex relations between the contents of episodes — thoughts, utterances, signals — and their other properties. In the reflexive-referential theory we advocate, episodes and their properties are not only kept in mind, but in the theory.

Acknowledgments The authors want to thank the Spanish Government (FFI2015-63719-P (MINECO/FEDER)) and the Basque Government (IT1032-16) for their grants, and the Institute for Logic, Cognition, Language and Information (ILCLI) of the University of the Basque Country (UPV-EHU), the Department of Philosophy and the Center for the Study of Language and Information (CSLI) at Stanford, and the Diamond XX Philosophy Institute in Copperopolis, CA, for the many direct and indirect benefits the three of them have had through their involvement with the activities of these institutions.

¹⁰Billy Joel "You might be right," in *Glass Houses* (1980).

References

- Barwise, J., & Perry, J. (1983). *Situations and attitudes*. Cambridge, MA: MIT Press.
- Crimmins, M., & Perry, J. (1989). The prince and the phone booth: Reporting puzzling beliefs. *Journal of Philosophy*, 86(12), 685–711.
- de Ponte, M., & Korta, K. (2017). New thoughts about old facts. On Prior's root canal. In M. de Ponte & K. Korta (Eds.), *Reference and representation and thought and language* (pp. 163–178). Oxford: Oxford University Press.
- Frege, G. (1879 [1967]). *Begriffsschrift, Eine der Arithmetischen Nachgebildete Formelsprache des Reinen Denkens*. Halle. (Translated by Stefan Bauer-Mengelberg, *Begriffsschrift, a Formula Language, Modeled upon that of Arithmetic, for Pure Thought*. In Jean van Heijenoort ed. 1967. *From Frege to Gödel: A Source Book in Mathematical Logic, 1879-1931*. Cambridge, MA: Harvard University Press).
- Kaplan, D. (1979). On the logic of demonstratives. *The Journal of Philosophical Logic*, 8, 81–98.
- Kaplan, D. (1989). Demonstratives. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan* (pp. 481–563). New York: Oxford University Press.
- Korta, K., & Perry, J. (2011). *Critical pragmatics*. Cambridge: Cambridge University Press.
- Lewis, D. K. (1979). Attitudes De Dicto and De Se. *Philosophical Review*, 88, 513–543.
- McTaggart, J. M. E. (1908). The unreality of time. *Mind*, 17, 457–473.
- Perry, J. (1977). Frege on demonstratives. *Philosophical Review*, 86(4), 474–497. Reprinted in Perry, 2000, Chapter 1.
- Perry, J. (1979). The problem of the essential indexical. *Nous*, 13(1), 3–21. Reprinted in Perry, 2000, Chapter 2.
- Perry, J. (1988). Cognitive significance and new theories of reference. *Noûs*, 22, 1–18. Reprinted in Perry, 2000, Chapter 11.
- Perry, J. (2000). *The problem of the essential indexical and other essays*. Stanford: CSLI Publications.
- Perry, J. (2003). Predelli's threatening note: Contexts, utterances, and tokens in the philosophy of language. *Journal of Pragmatics*, 35, 373–387.
- Perry, J. (2012[2001]). *Reference and reflexivity* (2nd ed.). Stanford: CSLI Publications.
- Predelli, S. (1998). I am not here now. *Analysis*, 58, 107–112.
- Predelli, S. (2011). I am still not here now. *Erkenntnis*, 74, 289–303.
- Prior, A. (1959). Thank goodness that's over. *Philosophy*, 34, 12–17.
- Wettstein, H. (1986). Has semantics rested on a mistake? *The Journal of Philosophy*, 83, 185–209.

Indirectness and Intentions in Metasemantics



Michael Glanzberg

Abstract This paper argues in favor of an indirect metasemantics for the standard for gradable adjectives. Specifically, it argues that multiple factors work to fix this parameter's value, and it is a further matter of context just how those multiple factors combine in any given case. The paper then asks how much an indirect metasemantics must depart from an intention-based metasemantics. The indirect metasemantics presented gives speakers' intentions a limited role, but more fully intention-based but indirect metasemantics are possible. The paper goes on to argue in favor of a less intentional metasemantics, via observation of the role of non-intentional aspects of cognition in fixing standards for gradable adjectives. The paper ends by considering implications of such a view for the nature of communication.

Keywords Metasemantics · Gradable adjectives · Intentions · Cognition

There are many context-dependent expressions in language. Obvious examples include indexicals and demonstratives like *it* or *that*, which require supplementation from the context to fix their referents. In recent years, there has been a great deal of discussion of other sorts of cases, where context dependence does not simply

Thanks to Jeff King for many discussions of the material in this paper and comments on an earlier draft. Thanks also to Josh Armstrong, Chris Barker, Ernie Lepore, Karen Lewis, Peter Pagin, Paul Pietroski, Jessica Rett, Lance Rips, Rob Stainton, Matthew Stone, and Alexis Wellwood for many helpful comments and discussions. Versions of the material in this paper were presented at the Rutgers Center for Cognitive Science Semantics Workshop, September 2012, and the Conference on Philosophy of Language and Linguistics at the Inter-University Center, Dubrovnik, September 2017. Thanks to all the participants at those events for helpful and lively discussions.

M. Glanzberg (✉)
Northwestern University, Evanston, IL, USA
e-mail: m-glanzberg@northwestern.edu

attach to an overtly pronounced word, but to some other linguistic component. Much attention has been paid to cases where the component is a hidden or unpronounced expression, that is only revealed through semantic or syntactic analysis. Many of these take the form of parameters in the appropriate syntactic representations of sentences, whose values are set by context.¹

It has proved fruitful to distinguish several different questions about these sorts of parameters. First, of course, what are they, and if they are not pronounced or are otherwise hidden, how do we know they are there? Second, what are their semantic properties; and specifically, what semantic values do they have? Third, how does context fix those values? This final question asks about what has come to be called the *metasemantics* of contextual parameters.²

There are two linked issues about the metasemantics of contextual parameters I shall explore in this paper. First, in earlier work (Glanzberg 2007, 2016), I have argued that different sorts of context-dependent expressions, including different sorts of contextual parameters, have different metasemantics. In particular, I distinguished what I call *direct* from *indirect* metasemantics. The model for direct semantics is that of demonstratives. It is especially clear on an intentional account of the metasemantics of demonstratives, where a speaker's referential intention fixes the referent of a demonstrative. In this case, one single factor determines its value. In contrast, I have argued that many contextual parameters have an indirect metasemantics, on which multiple factors work to fix the parameter's value. It can be a further matter of context just how those multiple factors combine in any given case. The example I shall focus on here, as I did in other work, is that of the *standard* for a gradable adjective. Take a simple case like:

- (1) Max is tall.

How tall someone needs to be to count as tall is a matter of context. The *standard*, as I shall discuss in more detail in a moment, is the parameter that determines this. It sets how tall you have to be in a given context to count as tall. I have argued that multiple factors influence how the value of this parameter is set in context,

¹The literature on this is large, but highlights include Bach (1994), Perry (1996), Recanati (2004), Sperber and Wilson (1998), Stanley (2000), and Stanley and Szabó (2000). I explore my own take on this issue in Glanzberg (2016).

²This distinction comes from discussions of reference, where Kaplan (1989a) and Stalnaker (1997) observed that it is useful to distinguish the semantic value of a referring expression from how that value is fixed. (Kaplan introduced the term 'metasemantic'.) The semantic/metasemantic distinction applies naturally to context-dependent expressions, where the question of how a semantic value is fixed thus becomes the question of how context fixes it. This is a part of the general area of pragmatics, but it is a highly focused part, so using the term 'metasemantics' to single it out has proved useful (even if it is jarring to some to see it called 'metasemantics'). The extension of the distinction to context-dependent expressions was mentioned by Kaplan (1989a). It was made explicit by Stanley and Szabó (2000) and discussed extensively in my (2007). Since then, the metasemantics of context-dependent expressions has become lively research area. See e.g. King (2014a).

including salient objects, speakers' intentions, lexical meaning, and so on. Thus, I have claimed, the standard has an indirect metasemantics.

The second issue I shall explore is the place of speakers' intentions in metasemantics. The model of direct metasemantics I have used is that of an intentional metasemantics for demonstratives, while the case I have made for indirect metasemantics highlights factors that go beyond speakers' intentions. This raises the question of how closely the two are linked. In particular, it raises the question of whether a metasemantics based solely on speakers' intentions must be direct.

My first goal in this paper is to review and expand my arguments in favor of indirect metasemantics, still using the standard as my example. My second goal is to ask how much an indirect metasemantics must depart from an intention-based metasemantics. The indirect metasemantics I shall present and defend includes much more than speakers' intentions, and gives such intentions a limited role. But I shall also note an option for a more fully intention-based but indirect metasemantics. I shall argue in favor of my own version. But I shall also note that deciding between them requires deciding a range of complex issues, and I shall not resolve all of them here. It requires deciding not only a range of specific issues about cognition, but some fundamental ones about how communication works. Thus, I shall make a case for indirect metasemantics, and a conditional case that it should not be thoroughly intentional, while highlight some important considerations that go into choosing which form a metasemantics should take.

I shall structure my defense of indirect metasemantics around a recent paper of King (2014a). King offers a carefully crafted intention-base metasemantics, and challenges my preferred metasemantics for the standard for gradable adjectives. I shall first review and strengthen my arguments for indirect metasemantics in Sect. 1. Next, I shall review and examine King's alternative in Sect. 2. I shall argue there that King must adopt an equally indirect, but intention-based alternative. I shall defend my preferred, less intentional version in Sect. 3. I shall also highlight there some difficult issues that could decide the choice between my approach and King's. I shall conclude briefly in Sect. 4.

1 Indirect Metasemantics

In this section, I shall make my case for indirect metasemantics, concentrating on the standard for gradable adjectives. The case will be supplemented in some ways in the next section, but the main set of arguments will be presented here.

Before getting to those arguments, I shall begin by reviewing some assumptions about semantics which will be our starting points.

As we will be discussing gradable adjectives at length, let us begin by reviewing their semantics. Gradable adjectives are those like *long*, *fast*, *bright*, *heavy*, *large*, *happy*, *cold*, *smart*, etc. They take degree modifiers like *very*. Roughly, they seem to report degrees of the relevant property. This is embodied in a common degree analysis of their meanings.

This analysis takes gradable adjectives to be functions from individuals to degrees on a scale.³ For example, the meaning of *tall* is given by a function to degrees on a scale (called a *measure function*):

$$(2) \llbracket \text{tall} \rrbracket (x) = d \text{ a degree of tallness}$$

This analysis is very natural for comparative constructions. For instance, we have⁴:

- (3) a. Max is taller than Mary.
 b. **tall(Max) > tall(Mary)**

(Of course, this is abstracting away from a lot of details about the comparative construction.)

The analysis of the positive form, i.e. the non-comparative form like *Max is tall*, makes it like an implicit comparative, comparing the degree assigned to an individual to a contextually provided value, called the *standard*. A simple version is:

- (4) a. Max is tall.
 b. **tall(Max) > d_c**

Here d_c is the standard, which is a contextually provided degree of tallness. To be tall in a context is to have degree of tallness greater than d_c .

To show just a little more compositional structure, note that we can give the comparative and positive forms related syntactic analyses along the lines of Kennedy (1997, 2007):

- (5) a. [_{DegP} [_{Deg'} [_{Deg} -er] [_A tall]] [_{PP} than Mary]]
 b. [_{DegP} [_{Deg} *pos*] [_A tall]]

As we are treating adjectives as measure functions, the semantics of *pos* is then:

$$(6) \llbracket [\text{Deg } pos] \rrbracket = \lambda g \lambda x. g(x) > d_c$$

Again, this is a rough-and-ready version of the analysis, but it is enough detail for our purposes.⁵

I shall return to some issues about the status of *pos* in Sect. 3. For now, we may observe that this brief glance at the semantics of gradable adjectives shows the presence of a contextual parameter: the standard d_c . It is implicit, as we do

³This theory has a long history. See among places Barker (2002), Bartsch and Vennemann (1973), Bierwisch (1989), Cresswell (1977), Heim (1985), Kennedy (1997, 2007), Rett (2015), and von Stechow (1984). Alternative theories have been developed by Burnett (2014) and Klein (1980).

⁴ $\llbracket \alpha \rrbracket$ is the semantic value of α . Were I need to mention specific semantic values, I shall put them in boldface. So, $\llbracket \text{tall} \rrbracket = \mathbf{tall}$.

⁵The references in Footnote 3 also discuss the positive form, and disagree over some of the details. For instance, Barker (2002) and Rett (2015) do not use this sort of *pos* analysis. The version presented here follows Kennedy (2007), though this is not his final proposal for the semantics of *pos*. For an extensive overview of the comparative, see Morzycki (2016).

not pronounce it. It does not come from nowhere, of course! On the analysis we are considering, it is contributed by *pos*, which is a silent Deg morpheme building the positive form. d_c is thus what in other work (Glanzberg 2016) I have called a *functional* parameter, as it comes from a functional element like Deg, and it behaves very differently from overt context-dependent elements like demonstratives.

Functional parameters, I have claimed, receive an indirect metasemantics. Focusing on the standard, I shall review and develop my case for this claim. To begin, let us start by considering some ideas about how d_c might be set.

A natural and common view of the way the standard is fixed is that it is derived from a contextually given comparison class via an appropriate norm. This is made vivid when we use *for*-phrases (cf. Klein 1980), like:

- (7) Max is tall for a professional basketball player.

This clearly pushes the standard up very high. Apparently, the average height of an NBA basketball player in 2016–17 is 6’7”.⁶ Assuming our norm is an average, we might expect the standard for (7) to be near to 6’7”, which is much higher than it would be in many contexts.

This view is part a semantic view, and part a metasemantic one. The semantics implicit in the comparison class view, and explicit in Bartsch and Vennemann (1973), is that:

- (8) $\llbracket [\text{Deg } pos] \rrbracket = \lambda g \lambda k_{(e,t)} \lambda x. g(x) > norm(k)(g)$

Here k is a comparison class, and $norm$ is function that determines the central tendency of g -degrees of members of k . For specificity, we can safely assume $norm$ is just the mean value of g -degrees of members of k .

Semantically, what before we presented as d_c is here given by $norm(k)(g)$. So, our parameter is really k , and its semantic value is a class of individuals. The metasemantics must fix the relevant comparison class, which is usually taken to be a class of salient individuals in the context.

This would be an option for a *direct* metasemantics. Though it is not that much like the metasemantics of demonstratives, it posits one specific contextual factor—salient individuals—to fix the value of the parameter k .

In earlier work (Glanzberg 2007), I have followed Kennedy (2007) in rejecting this view. The main reason I reject it is that there are contexts where the norm value on the comparison class is not the standard.⁷ Kennedy observes that the analysis that fixes the standard via a norm of being average applied to a comparison class (expressed by a *for*-phrase) predicts that the following should be contradictory:

- (9) Nadia’s height is greater than the average height for a gymnast, but she is still not tall for a gymnast.

⁶According to <http://www.basketballinsiders.com/2016-17-nba-roster-survey/>.

⁷Kennedy also argues that *for*-phrases trigger a *presupposition* that objects fall within the class fixed by the *for*-NP, rather than providing a comparison class argument.

It is easy to find contexts where this is not contradictory. We can also find contexts where it is not contradictory if we change the norm, e.g. to the mode or median height, and likewise for any other norm that is plausible given the meaning of *tall*. Hence, no fixed norm applied to a comparison class will be able to uniformly get the right truth conditions.⁸

With Kennedy, I concluded that the standard is not uniformly computed from a comparison class, and that a comparison class is not itself a contextual parameter in a gradable predicate. The comparison class is certainly a factor in determining the standard, and as we will see in a moment, an important one. But it is not semantically coded to do the job all by itself, and the metaseantics must take into account more than comparison classes.

Following Kennedy and McNally (2005), we can observe another important factor in some cases. The scale structure that an adjective lexically encodes can be crucial. Kennedy and McNally observe that ‘absolute’ adjectives like *open* and *straight*, which have a minimum or maximum value for their scale, at least usually require the extreme value to be taken. Consider:

- (10) a. The door is open.
b. The rod is straight.

To be straight is to be completely straight, and so to have the maximum value on the scale. To be open is to be even a small amount open, and so to have above the minimum value on the scale.⁹ Crucially, we see little or no context dependence in absolute cases. The reason is that the lexical meaning of the adjective does all the work, and nothing else about the context (typically) matters. Hence, comparison classes can be overruled by lexical meaning.

We thus have at least two factors that contribute to setting the standard: comparison class, and lexical meaning. Both can contribute, but neither is sufficient to fix standards for all adjectives in all contexts.

I have argued that other factors can play a role too. Most of the kinds of things that influence context can. Not surprisingly, shared presuppositions can (as we learned from Stalnaker (1978), among many others). Consider:

- (11) Context: We are at a summit on international development. Shared presuppositions about the problems facing poor nations, and the things that can be done to fix them, make it clear that individuals or countries can only make an impact with contributions of \$10 million or more. I make manifest my intention to count average philosophers as rich, by holding up photos of conferences and saying ‘look at all those rich people’.

⁸Kennedy (2007) and Fara (2000) also discuss ways that the norm plus comparison class analysis does not adequately address problems related to vagueness and the Sorites paradox. Kennedy credits the observation in (9) to Bogusławski (1975).

⁹Kennedy (2007) and Kennedy and McNally (2005) argue that these must take the limit value. A more modest position is that they usually do. Related ideas are discussed by Cruse (1986) and Winter and Rotstein (2004).

Assertion: We should try to get money from philosophers because they are rich.

Absent any other clear contextual information, the claim that philosophers are rich just appears false in this context (though of course, it could be true in a context where we were talking about, say, relative income in the United States). In this case, the shared presuppositions about our purposes in classifying as rich beat out an attempt to make a salient comparison class, and they beat out a fairly clear communicative intention.

In other cases, a comparison class can exert a stronger influence than common ground information that is presupposed by all parties. For instance:

- (12) Context: We surrounded by basketball players, and everyone is talking about basketball. I happen to be obsessed with jockeys, and you know that I am. A 5' 5" high man walks into the room. I say:

Assertion: He is tall.

It appears very hard to hear this as true. In this case, the standard is influenced by the salient group in the immediate environment, which can provide a comparison class. Our shared knowledge of my obsession with jockeys, and the communicative intention that goes with it, does not overrule this.

Work in pragmatics has also shown how previous discourse can affect context in different ways than mere presupposition.¹⁰ We see that at work with standards too. If we supplement (12) with some discourse, the situation changes:

- (13) Context: We surrounded by basketball players, and everyone is talking about basketball. I happen to be obsessed with jockeys, and you know that I am. A 5' 5" high man walks into the room.

Discourse: I am going to a horse race later today. I just know that the shorter jockey will win. They always do. You have to be under 5' to do really well.

Assertion: He is too tall.

I am not sure how firm the judgement is here, but it is much more easy to hear as true. So, overt discourse can override prior presuppositions, and can implicitly guide us to different comparison classes we well.

We have thus seen a number of factors that can contribute to setting the value of the standard parameter in context. Comparison classes, standing presuppositions, overt discourse, and lexical meaning, can all play a role. There are presumably more factors that can, but these are some. Depending on the context, we can find that different elements play a more or less central role. This makes the metasemantics of the standard indirect.

The point of the indirect metasemantics is that these competing features have to be combined, and combining them is itself something that takes place in context.

¹⁰Some classic examples of this are from work of Heim (1992) and Roberts (1996).

There are general rules that might be invoked. Kennedy suggests an interpretive economy principle that asks us to maximize the use of lexical content over contextual factors. But like most contextual rules, this is defeasible. A contextual parameter with an indirect metasemantics must be set by the various pieces of information context provides, but context does not simply hand us a value for such a parameter, nor does it hand us a uniform rule for computing the value from a specific piece of contextual information. Rather, a range of contextual information and computational rules must be taken into account and weighed in working out the value from context.

Finally, we should ask if experimental work can help shed light on what metasemantics the standard must have. Not that much is known about how we compute standards in real time. However, I do believe that what experimental evidence there is compatible with the indirect metasemantics I have advocated.

Experiments conducted by Barner and Snedeker (2008) show that even by age 4, children are able to use statistical properties of objects in their environment to compute standards, and also use clues from nouns to select comparison class objects. In a somewhat related vein Schmidt et al. (2009) also found that adults performed some statistics on comparison classes. They compared several models of what those statistics might be, and got results preferring either a range-based model (the top percentage of the range) or a clustering model developed for general categorization. Both of these were preferred over a model based on mean and standard deviation. Solt and Gotzner (2012) also confirm that adults are sensitive to the statistics of comparison classes. They have somewhat different findings than Schmidt et al. (2009), and in particular highlight the role of rich measurement systems, allowing some way of defining distance between degrees. (Hence, when it comes to the pragmatics of the standard, an interval scale is not sufficient.) They show this result both for adjectives associated with numerical measurement systems and those without.

Syrett et al. (2010) focused on the difference between relative and absolute adjectives, and found evidence that both children and adults assign substantially context-invariant meanings to absolute adjectives, but are sensitive to contextual standards for relative ones. Rips and Turnbull (1980) find a similar contrast between relative gradable and non-gradable adjectives in adults. Studying adults using a variety of methods (response times for acceptability judgments, eye tracking in reading, choice of interpretation, and acceptability rating), Frazier et al. (2008) found that the difference between maximum and minimum standard adjectives proposed by Kennedy and McNally (2005) and Winter and Rotstein (2004) is respected in comprehension, and that the scale structure is accessed in the course of comprehension. They also find support for Kennedy's interpretive economy proposal, and add a kind of locality in processing constraint to supplement it.

I doubt these results conclusively show what happens in the real-time processing of standards, but they are compatible with the indirect approach. They show clear roles for different factors, including comparison classes and lexical meanings. This is a key part of the indirect view.

These results also highlight something else that is important about the indirect view. The experimental results show how we, even as children, are able to implicitly do substantial statistics in the course of our pragmatic interactions. This is not an isolated phenomenon. We are simply good at implicitly doing some statistics.¹¹ But we do it in a highly sub-personal way, constituting more a sensitivity to statistical facts in our environment than an overt intention to compute such statistics. Indeed, the algorithms we use might be highly complex. (For instance, one of the models Schmidt et al. (2009) consider works with partitions on the comparison class, and Gaussian distributions defined over those partitions.) But the complexity is not by itself a problem. It is an interesting fact about us that we are able to do such complex statistics, even as children. Complex or not, the evidence is that we can. But it is a distinctive ability that we bring to a number of learning and reasoning tasks, and one that is highly tacit. As I shall discuss more in Sect. 3, this kind of distinctive and tacit ability marks another important aspect of the indirect approach. It is a distinct cognitive ability we can access in fixing standards. It can make a contribution to fixing standards, on the indirect view.

We thus have a case for the indirect metasemantics of the standard. We have seen multiple factors that can contribute to fixing the standard, including comparison classes, presuppositions, overt discourse, lexical meaning, and presumably more. In different contexts, these are weighed differently. In context, we must combine the various factors, by whatever defeasible rules we might rely on. In doing so, we can access cognitive abilities, like our ability to do tacit complex statistics on comparison classes. Thus, we have a metasemantics that is very indirect.

2 The Role of Intentions

The version of indirect metasemantics I have just articulated and defended downplays the role of speakers' intentions. They are certainly a factor that can contribute, but they are one among many. As I presented them, comparison classes in the immediate environment of a context and lexical meaning are factors that are not understood as speakers' intentions. Indeed, we saw cases above where these sorts of factors can override a speaker's communicative intentions.

In earlier work (Glanzberg 2007, 2016), I argued explicitly that there is only a limited role for intentions in this metasemantics. I shall now reconsider that argument, and examine a challenge to it from King (2014a). King argues in favor of a uniform intentional analysis of the metasemantics of context-dependent expressions, and argues I underestimated the role intentions play in metasemantics. My main goal in this section is to support my argument for indirect metasemantics,

¹¹This is a very broad idea in psychology, with a huge literature. Some references, specific to language learning, are Frank et al. (2009), McMurray et al. (2012), Smith et al. (2014), Xu and Tenenbaum (2007), and Yu and Smith (2012).

by showing that King's favored version is still, I claim, indirect. But in doing so, I also want to reconsider where and how intentions might play a role. In the next section, I shall argue in favor of my less thoroughly intentional approach.

In addition to the importance of what I see as non-intentional factors, the main reason I argued for a limited role for intentions is that I doubt speakers will often have intentions that are specifically about the standard. They will often not have any idea what such a thing is, which limits the sorts of intentions they can have. To make matters worse, the standard comes from *pos*, and it is doubtful many speakers have any idea there is such a tacit piece of syntax towards which to direct intentions. This is not to say speakers have no intentions that relate to classifying or comparing that affect the standard value. They usually will have many of those intentions. But they will often not be intentions to refer to or otherwise pick out a single standard value, and they will only in very rare cases be intentions to set a value of a parameter provided by *pos*. This is in contrast to an intentional direct metasemantics for demonstratives. In that case, speakers can have full and overt referential intentions to pick out a specific object by using a demonstrative, and intend that object to be the referent of the demonstrative expression.

We see speakers' lack of such specific intentions about the standard in how difficult it is for speakers to answer queries about the standard in many cases, and how hard the queries can be to formulate successfully:

- (14) a. Max is rich.
 b. i. How rich is rich?
 ii. What do you mean rich?
 iii. ? By what standard?
 c. i. Well, you know, really rich.
 ii. Well, you said Mary was rich, and so is Max.
 iii. We are talking about Google.

This contrasts with cases of demonstratives, where we can ask about referential intentions and in most cases get clear answers:

- (15) Context: Looking at two paintings right next to each-other in a museum.
 a. That is beautiful.
 b. i. Which one do you mean?
 ii. Which one are you intending to refer to?
 c. i. The one on the left.
 ii. I was talking about the one on the left.

With demonstratives and the referential intentions that back them, we can query directly about the referent and get a clear answer. With the standard, we often get only partial or indirect answers, which do not seem to refer to the standard in many cases. Similarly, we cannot often query about the standard directly, but have to come up with indirect ways to ask about what fixes it. The demonstrative case shows

evidence of a direct metasemantics, probably an intention-based one, at work. The standard case shows evidence of an indirect metasemantics making only minimal use of intentions.¹²

There are speaker intentions here, of course, but not like referential ones. They are better described as intentions to compare and measure. Indeed, we saw how some of intentions can work to fix standards in examples like (11) and (12). But we also saw that those sorts of intentions are among many competing factors. From these intentions, plus the effects of the wider context, including previous discourse and the environment in which a claim is being made, a standard emerges. As the role of comparison classes makes vivid, often salient examples or groups have a great deal to do with how the standard is fixed. But there is often nothing like a single referential intention at work.

The indirect view thus makes room for intentions, but they play a limited role, as part of a broader indirect metasemantics. They need not be the determining factor, and may not be available in all contexts.

In recent work, King (2014a) has challenged this conclusion. He argues that I underestimated the role of intentions. I shall discuss King's own view, and how it relates to mine, in a moment. But first, I want to make a concession on this issue. My position in Glanzberg (2007) was that we cannot have intentions specifically about the standard. I just ran through some arguments for this claim, but I put it more cautiously that we often do not. King pointed out ways that we can sometimes have intentions directed towards a standard, in spite of these arguments. I have implicitly already conceded to King that we can indeed sometimes have such intentions. Hence, I have put my claim now that we often do not. To make this concession clear, I now hold that this is as strong a claim as the evidence justifies.

In fact, it appears that we can sometimes have relatively clear intentions to set standards. Here is an example. On the internet, you may find a group of *Tall Clubs International*.¹³ They list the following rules:

- (16) TCI Membership requirements. Minimum height: 6'2" for men, 5'10" for women, or taller, when measured in stocking feet.

In this kind of case, it seems we do have an intention that fixes the standard. We may not call it 'the standard', but we intend to fix a cut-off for what counts as tall. It seems we can have intentions to set standards, at least in some cases. In some cases, these can play a key role in setting the standard.¹⁴

¹²I have adopted an intention-based direct metasemantics for demonstratives as a model of a direct metasemantics. The classic issue for the metasemantics of demonstratives is whether it is speakers' intentions or something overt, like a gesture, that fixes their reference. Famously, Kaplan (1989b) opted for the latter, but then changed his mind in Kaplan (1989a). I have not argued for the intentional alternative, but King (2013, 2014b) develops a good case for a sophisticated intention-based view.

¹³<http://www.tall.org/>. They have member clubs world-wide.

¹⁴One more minor point. In earlier work (e.g. Glanzberg 2007), I relied on a formulation of the semantics of gradable predicates that closely followed Kennedy (2007). This version relies not

Even so, there are several questions this raises. One is how often this sort of case happens. We may have intentions in some special cases about height or wealth, but do we about brightness, roughness, etc.? There is work that tries to quantify brightness in psychophysics,¹⁵ and perhaps if you know enough about it, you could intend to set a value for brightness the same way the Tall Clubs International intend to set a standard for tallness. But this seems to be a highly unusual case, and not the normal way speakers proceed. Moreover, the kinds of examples I gave in (14) make me doubt that substantial intentions directed towards standards are that common, even in cases where we *can* form such intentions. In many cases, we operate without such specific intentions. The contrast with (15) also makes me suspect that when we do have such intentions, they are often not like the referential intentions involved with demonstratives. Finally, even if we have intentions to set standards, these may not go with any grasp of what parameter is being set, as speakers may still not know about *pos* or *d_c*. Unlike referential intentions, these intentions still seem more about how to classify, and not how to fix the value of an expression we use.

I do not challenge that in most any utterance, there are a wealth of communicative intentions at play, and the indirect metasemantics makes room for them. But we often do not have intentions specifically about standards. When we do have such intentions, they are often not like the ones we use to pick out an object and publicly refer to it. I concede that intentions can do those things, but often they do not. The indirect metasemantics allows for both options.

With that in mind, I want to turn to King's positive proposal. In a series of works (King 2013, 2014a, b), King has advanced a uniform metasemantics for virtually all context-dependent expressions (of a class he calls 'supplementives'). King's account is much more thoroughly intentional than the one I have offered for the standard for gradable adjectives (though I do accept his account for demonstratives). Even with the concession I made above, his is a substantially different view. Perhaps more importantly, given that his metasemantics is the same for demonstratives and the standard, we might wonder if he thereby defends a direct metasemantics for the standard. I shall now examine how direct King's metasemantics is. I shall argue that for the standard, he winds up with a position that is in fact *indirect*, though intention-based. This in the end supports my claim that some metasemantics is indirect, as a leading alternative to my preferred view is also indirect. Once that is established, we can return to the issue of how intentional a metasemantics should be. This will be the focus of the discussion to follow in Sect. 3.

King's positive view is that he calls the 'coordination account'. In outline, the account of this (King 2014a, p. 102, with some notational changes):

on a standard value *d_c*, but on a function *s* that takes lexical and contextual inputs. I was mostly concerned to follow Kennedy in making vivid the role of lexical content in fixing some standards, but I have since come to realize that Kennedy's proposal is actually stronger than just that. King (2014a) points out that it may be easier for speakers to have intentions about standard values *d_c* than about this functions. For this reason, and to avoid some of Kennedy's specific commitments, I now prefer the *d_c* version.

¹⁵A classic in this literature is Stevens (1975).

- (17) The semantic value of a context-dependent expression e in a context c is that element (object, or appropriate other value) o that meets the following two conditions:
- a. The speaker intends o to be the value of e in c .
 - b. A competent, attentive, reasonable hearer who knows the common ground of the conversation at the time of utterance would know that the speaker intends o to be the value of e in c .

In short, the speaker intends the term to have a value, and the intention must be known to anyone who knows the common ground. The intention must thus be manifest in the conversation somehow.

This is clearly an intention-based metasemantics. It also appears like it might be uniformly direct. Not only does it give demonstratives and the standard the same metasemantics, it makes that metasemantics one of what appear to be referential intentions; at least, intentions to set semantic values. King's version is a highly sophisticated intention-based metasemantics, and I shall not challenge it as a metasemantics for demonstratives.¹⁶ I am willing to extend it beyond overt demonstratives, to what I have called *thematic parameters* (Glanzberg 2016). But the question remains whether it really offers a direct metasemantics for functional parameters like the standard, where I maintain we need an indirect metasemantics of the kind I discussed above.

King's idea for the standard (King 2014a) is that the speaker's intentions may determine a cut-off degree in a different way than we might expect in simple cases of demonstratives.¹⁷ He gives the example of *smart*. We may not have intentions about degrees of smartness, but we might have intentions about *kinds* of people who will count as smart. Presumably even if we do not know about the parameter d_c or *pos* or degrees of smartness, we do know that it is gradable and we need to set a cut-off for *smart*. We do that via our intentions towards kinds. We direct our intentions towards kinds of people and those in turn determine a standard for *smart*. More generally, we can direct our intentions towards kinds of objects, and then those in turn determine an appropriate standard.

In the end, I shall ask if it is most useful to think of what happens in cases like King's as only involving intentions. But first, I claim that King's metasemantics remains *indirect*, in spite of initial appearances.

As always, we should start by contrasting the case of demonstratives with that of the standard. For a perceptual demonstrative, for instance, we might expect the speaker to perceptually individuate an object, intend to refer to it, and make that intention public via pointing or some other way of manifesting a referential

¹⁶My own view is that in some special cases, the speaker's intention can fix the value of demonstrative even if it is not fully manifest. So, I do quibble with some details of King's analysis. But it is not my goal to argue this point here, and King's developments do offer a very strong form of the intentional view of the metasemantics of demonstratives.

¹⁷As I discussed above, maybe in some special cases we have something more like ordinary referential intentions.

intention. They will know which expression's value they thereby fix. For the standard, on King's theory, the speaker will intend to count a certain kind of people as smart. That in turn fixes a standard for smartness. They may understand that this creates a cut-off, but they may well not know there is a parameter whose value needs to be set. King's metasemantics thus marks a clear contrast between simple cases for demonstratives and cases like the standard. And for the standard, the metasemantic process King proposes already has some indirectness, as it runs through a kind to reach a degree value.

Many of the points I made above in favor of an indirect metasemantics can be re-cast in King's intentional terms. This will make clear just how indirect King's proposal must be in the long run, if it is to handle a full range of cases. First, let us look at the role of kinds of individuals. King is right that we can sometimes fix standards via kinds. But the evidence, both from examples like (12) and the experimental evidence I reviewed in Sect. 1, shows that we can also fix standards by reference to statistical properties of salient individuals in a comparison class. We might have intentions corresponding to those salient individuals, which we might express in terms like *the salient things around here*. So we have at least two comparison class-oriented options for how our intentions might go: intention to fix a standard via a kind, or via the statistical regularities across salient individuals. We have multiple options, and both make an indirect path from what we intend to the resulting value.

We have also seen, both from examples like (10) and experimental evidence, that lexical content can change the situation, leading to much less context dependence, and little reliance on either of the options we just considered. If we are to take King's line, that would relate to a different sort of intentions, specific to the lexical meanings of absolute adjectives. We might, for instance, have a tacit intention to describe things as *completely open*. So, we have multiple roles for comparison classes and a role for lexical content as part of our metasemantics, even if we view it as intentional. Each comes with different kinds of intentions, and the intentions can be to varying degrees complex and tacit.

The same can be said for the role of prior discourse. Now, we might see prior discourse as making parts of our complex communicative intentions manifest. And so, we might describe the affects of discourse on the standard in intentional terms. But still, it changes the options for how we are to integrate salient individuals into our process of determining the standard, as we see with example (13). We also see with example (11) and its contrast with (13) that presuppositions can compete with overt discourse for fixing the standard. If these are all intentions, or the making of intentions manifest, then we have multiple kinds of intentions or ways to make them manifest. In different contexts, these contribute differently to fixing the standard, and how they contribute can be a matter of context. This is a re-casting in intentional terms of the full case I made above for an indirect metasemantics.

What this shows, I believe, is that if we wish, we can offer an *intention-based indirect metasemantics* for parameters like the standard. We can describe what

fixes the standard as a complex set of communicative intentions, making reference to multiple sorts of contextual factors, that must be integrated in a contextually appropriate way to indirectly fix a standard. We can, with King, require those intentions to be publicly manifest. The contrast with demonstratives still stands. What we have is an indirect metasemantics. In light of the arguments I gave in Sect. 1 and here, I believe that the best way to understand King's proposal is precisely as an intention-based indirect metasemantics, when it comes to parameters like the standard.

I take this as further support for an indirect metasemantics for parameters like the standard. In addition to the arguments I gave in Sect. 1, we find that King's very different approach also winds up providing an indirect metasemantics when it comes to the standard.

But that leaves open the question of whether we should opt for King's intentional version, or the version I described above, which leaves only a limited role for intentions (even with the concession I made that there is a greater role sometimes!). In the next section, I shall argue that a fully intentional description of the metasemantics is not the best option. Rather, I shall claim, my original version that makes a place for intentions among many factors is a better description of how our metasemantics works. But I shall also note some issues that need to be decided before a final decision on the matter can be made.

3 Cognition and Semantics

Much of what happens in communication is intentional, and so the intention-based metasemantics is tempting. But I think there are reasons not to uniformly describe things this way. Not everything involved in fixing the standard is best seen through the lens of intentions.

To see this, let us consider how our communicative intentions interact with linguistic elements, both overt and hidden. To begin, let us consider this idea quite generally, rather than with regard to the standard in particular. Our communicative intentions are typically to describe events or states ('eventualities' of various kinds). We do this in language mainly by using a range of predicates and terms that capture important features of eventualities. Grammar provides us a range of lexical categories that express these. Verbs, nouns, and adjectives (and adverbs) provide us a rich range of predicates, and we build terms out of nouns, pronouns, and so on. The work of capturing our views of the world and communicating them is carried mostly by the lexical categories, and that is where our communicative intentions are mostly encoded and expressed.

As is well-known, language goes beyond this in providing so-called functional expressions, that provide what is sometimes loosely described as 'grammatical glue' that helps form sentences. Determiners, tenses, and so on are functional categories. It should be stressed that many of these have content, but they are not themselves

direct components of event descriptions; rather, they are part of how we structure those in our languages.¹⁸

To illustrate the difference in how we think and intend to communicate with the two sorts of categories, consider tense and time. We certainly often intend to place descriptions of events in time. But tense is not simply that. First of all, tense structures a clause syntactically. I doubt our intentions to place events in time relate that closely to the syntactic form of a TP (Tense Phrase), or the grammar of inflection. And second, tense expresses time in distinctive ways. Overtly, English has a past tense *-ed* (plus irregular forms), but has little if any marking of present tense, and future mostly comes out with the expression *will* that behaves more like a modal. And of course, we have many other kinds of temporal expressions, including adverbials and indexicals. Our intentions to place events in time intersect with the functions of tense in only a few ways. This is all the more so if our intentions go only with the overt structure of tense, as then we will often place events in time without using tense at all. If tense is tacitly present more than we see overtly in English, then still, our intentions do not clearly correspond with tense, as opposed to other overt forms or apparent surface forms. The functional ‘glue’ that builds clauses overlaps with our communicative intentions in some ways, but is not simply an encoding of them.

Now, we can return to our main topic of gradable adjectives and the standard. We will see there the same contrast between lexical and functional, and this in turn will help us to understand the role of degrees in our thinking and our intentions. Let us start with the major lexical category A (Adjective). Within that domain, the semantic hypothesis we are exploring is that gradable adjectives fundamentally express some form of measurement. We should expect, with the lexical category, that this fits with our ways of describing and communicating about eventualities. It should relate to communicative intentions in some close way. I shall return to this in a moment.

But as with all expressions, there is more to the story than that. Our grammars impose structure that fits concepts into the specific semantic and syntactic categories we use. (In other work (Glanzberg 2014) I have called this ‘packaging’.) Lexically, gradable adjectives have two linked features that make this happen. Their semantics provides dense linear ordered scales, and they need an extended projection DegP introducing the functional element Deg in order to compose in most cases.¹⁹ The latter is a common view in current syntax: most every lexical category first builds an ‘extended projection’ with higher functional elements.

Most of our evidence for both comes from the grammar of comparatives. If we ask why scales for adjectives are connected and transitive, we look to the comparative:

¹⁸See any syntax textbook, and the general discussions of Baker (2003), Fukui (2001), and Grimshaw (2005). Much of the original work on functional categories stems from Abney (1987). I have discussed some aspects of this in my Glanzberg (2008, 2014).

¹⁹As I mentioned in Sect. 1, there is some dispute about this, but it is a well-supported hypothesis, and it renders the points I shall make here most explicit. Other options for the syntax-semantic interface would get similar results in the end.

- (18) a. For any a and b of type D_e (restricted to material objects), a is heavier than b or b is heavier than a or a and b are equally heavy.
 b. If a is heavier than b , and b is heavier than c , then a is heavier than c .

Thus, I claim, it is the grammar of Deg, most visible in the comparative, that reveals the semantic structure by which we encode adjectival meaning. In turn, the semantic structure of adjectives is formatted to be able to project into Deg, to produce usable forms. But the presence of Deg is not transparent to many speakers, nor is the structure of overt comparative phrases. Nor is the packaging into dense linear orderings that they require. But that is the key piece of grammar that takes our general ability to measure and builds real lexical items. Thus, in the abstract sense of linguistic competence, we may ‘know’ these features of adjectival phrases are there (or behave in ways they dictate), but that comes from whatever gives us competence with abstract, functional aspects of grammar. We do not often voice specific intentions around those aspects of language.

I highlight these points, because they remind us that apparatus like precise scale structures and standards are products of grammar, whose relation to how we think and communicate itself can be highly indirect, as it is with tense. Humans, along with many animals including rats, are equipped with the ability to represent a range of magnitudes (e.g. Cantlon et al. 2009; Feigenson 2007; Meck and Church 1983). Well-studied ones include length, time, and pitch, and also brightness, warmth, weight, etc. Of course, the most-studied one is number (e.g. Carey 2009; Dehaene 2011; Gallistel and Gelman 1992). Many of these cognitive abilities correspond to adjectives like *fast*, *large*, or *bright*. On the basis of this, it is safe to assume that our measuring talk is linked to our abilities to measure various quantities, and to convey those measurements. Our adjectival meanings do reflect ways we can think about eventualities.

These representations are ‘analog’ or ‘approximate’, in that they give continuous representations even when the underlying phenomena are discreet. Their behavioral signature is Weber’s law: discrimination of magnitudes is a function of their ratios. For this reason, the relation between the ways we think about magnitudes and the ways they are semantically represented are themselves fairly indirect. This is much as we have seen with other cases, like the relation between our thinking about time and the grammar of tense. My own view is that for approximate magnitudes, our behavioral responses at best form a constraint on underspecified semantic representations. I do not need to argue that case fully here, but only to note that the semantics is not a direct encoding of our cognitive abilities.²⁰

Though our main topic is the metasemantics of the standard, this detour through the semantics, grammar, and cognition of adjectives makes two points about underlying indirectness for gradable adjectives, even before we look at the standard itself. There are at least two ways that the semantics and syntax of gradable adjectives already sets up indirectness. One, as we just saw, is that the mapping from

²⁰I argue this at length in work in preparation, and touch on it in Glanzberg (2014).

our magnitude representation systems to lexical representations is complicated. But also, our grip on the linguistic locus of magnitude and comparison is tacit. It comes from our tacit grasp of the structure of DegP. We have that, but we have in whatever way we tacitly grasp syntax, not in virtue of thinking about magnitudes. We thus find two features of indirectness in the syntax and semantics gradable adjectives:

1. Our ability to think about magnitudes and compare them is only indirectly mapped to lexical representations.
2. Our grasp of the linguistic encoding of magnitude and comparison is provided by our tacit grasp of the grammar of DegP, not our ability to measure and compare.

The link between our thinking and our intentions to communicate those thoughts and the forms of expressions we use in language and their meanings is already indirect, and involves multiple distinct cognitive capacities.

With all this in mind, let us finally return to the task of fixing the standard in context. Let me offer what I think the best description of what we do in many cases of that task. First, we start with our understanding of events and states, which provides us with thoughts we intend to communicate. When it comes to the properties expressed by adjectives, our appreciation of events often derives from our abilities to approximately represent magnitudes. Grammar already puts two steps between those thoughts and intentions and the fixing of a standard. It imposes a structure of precise degrees where we often have little precise thought. And, it structures the task as one of fixing a value for a functional parameter supplied by a functional head like *pos* that we only grasp through tacit knowledge of grammar.

Might those two suffice to fix a standard? In some cases, perhaps. I think the most likely cases are not those where we rely on approximate representations in our thinking, but where we rely on more overtly articulated concepts, often scientific ones. In cases like this, we might find the kind of specific intentions that we saw in example (16). We also might see this in mature uses of claims like:

(19) Styrofoam is not dense.

In cases like these, we might possess knowledge (often scientific knowledge) that could allow us to fix precise degrees on a scale. Using this, we could form intentions to fix cut-offs.

But even in these cases, where we might possess mature scientific knowledge that could allow us to fix a precise degree on a scale, we rarely have those sorts of intentions overtly. Why not? As I have mentioned, such intentions would involve a range of tacit and scientific knowledge that may or may not be available to speakers. But also, such intentions are often not necessary or even useful. The role of comparison classes is central here. We are in fact sensitive to statistics of comparison classes in our environment, and our communicative practices often rely on that. Grammar requires there is a precise value for implicit comparison, and often, it is our ability to respond to statistical regularities that steps in to make it happen. If we happen have the ability to identify the value explicitly in some cases (I think rare cases!) it will often not matter, as the grammar and statistics of implicit comparison can step in and do the job for us anyway.

So, we have at least three factors that go into our fixing a standard: grammatical competence, ability to represent the kinds of magnitudes in question, often only approximately, and our ability to respond to statistics.

But the moral of the arguments for indirect metasemantics is that there is more to the story even than this. Lexical meaning can override sensitivity to statistics. Overt discourse can change comparison classes, and override background presuppositions. Generally, we possess a number of ways to affect how parameters are set in rich conversational settings, and we can use any of them.

Communicative intentions are most vivid here with the roles of presupposition and overt discourse which can make a speaker's intentions clear. But intentions only function in combination with other aspects of cognition. Grammatical competence is, presumably, a highly special purpose ability (a module, I presume). Our ability to do statistics is a highly general-purpose ability we have that seems to affect a range of leaning and reasoning situations, but is likewise highly tacit, and not something we form overt intentions about. Our ability to represent magnitudes is partly shared with rats, which show little if any communicative intentions.

Now, if we stretch the notion of intention enough, I suppose that is alright to call any of these intentions. And in cases like (19), perhaps overt intentions can sometimes seem more important than the tacit knowledge of grammar that also plays a role. But always putting things in terms of a stretched notion of intention elides important distinctions. Some aspects of the metasemantics of the standard are familiar ordinary communicative intentions. But some are highly tacit abilities. Some are special purpose, while some are general-purpose. These can serve our communicative intentions without being intentions. That is clear from tacit knowledge of syntax, but it holds for tacit abilities to do statistics too. Fixing a standard is a process that draws on many of these distinct kinds of abilities. The version of the indirect metasemantics I prefer keeps track of these distinctions. It can appeal to intentions specifically, in cases like previous discourse, while appealing just as specifically to our abilities to do statistics or our grammatical competence. I think this is the most revealing way to capture the metasemantics of the standard. Hence, I do not opt for a thoroughly intention-based metasemantics. I thus conclude that the metasemantics of parameters is indirect, and also, that the indirectness is not best described in uniformly intentional terms. Intentions are crucial, but not the only factor.

There is a reply from defenders of intention-based metasemantics that we should pause to consider. I have highlighted a variety of cognitive resources we draw on to think about properties expressed by gradable predicates, to structure those thoughts in language, and to fix standards in context. Not all of these are intentions, and not all of them are fundamentally about communication. But still, we can employ basic, often tacit, cognitive resources to serve a variety of functions. Our perceptual abilities serve our motor abilities when we run to catch a ball, for instance. Is what we see with the standard genuinely a mix of intentions and non-intentional cognitive factors, or rather a mix of intentions, some of which are based on underlying tacit non-intentional factors? I have suggested the former gives us a more accurate picture

of what happens when we fix the standard, and so is a better metaseantics. But it is open to intentionalists to pursue the latter.²¹

I have already given my reasons why I prefer the less intentional version. It offers a more refined account of the facts, that keeps track of more information about how the standard is set. But there is a reason that intentionalists can offer for preferring their option. One might hold that if any of the non-intentional resources I have mentioned—tacit statistics, magnitude representation systems, grammatical competence, etc.—are to serve communication, they must do so via communicative intentions. If one holds this view, then the only role for these kinds of resources is to produce complex communicative intentions. If so, then the intentionalist response we just considered is obligatory, and nearly a tautology.

I do *not* hold this view. I shall not here attempt a full defense of my position. It is too large an issue to address in this paper. For the most part, I shall rest with the my defense of indirect metaseantics, and the conditional claim that if one does not take this view of communication, then my less intentional version of the metaseantics is better. But, I shall briefly discuss some points surrounding the two competing views of communication, and give some indication of why I prefer a less intentional view of it too.

As I have stressed, I do see a role for intentions in many places, including communication, of course! This is absolutely clear in the many cases of Gricean reasoning we see in communication. As I also said, I think when it comes to the main messages we wish to convey, there are communicative intentions at work.

But I also think that the scaffolding upon which we build our overt thoughts and communicative intentions need not be intentional at all, nor need the cognitive and linguistic abilities that provide the scaffolding underlie any specific communicative intentions. To put it simply, when it comes to lower-level cognitive and linguistic abilities, overlap of ability is enough, and we need not have intentions, shared intentions, or any other Gricean-inspired apparatus of communicative intentions corresponding to these. Such overlap of lower-level abilities is enough to serve communication and communicative intentions, even if it is not itself intentional.

Let us look at two cases that are easier than the standard: sound and syntax. We have substantial cognitive abilities to process speech sounds, which are known to depart from our abilities to hear sounds more generally.²² This ability behaves a lot like a module in the sense of Fodor (1983). Among things, it is domain specific and automatic. We cannot turn it off. At the same time, we as ordinary speakers have no idea how that mechanism works. It is extremely complicated. It is doubtful we have any intentions at all surrounding this ability. Indeed, it is hard to know what such intentions would be, short of the loose description of ‘wanting to say something’. We do not intend our speech sounds to take the shape they do, and we

²¹Though he has not endorsed this particular formulation, King (p.c.) has told me he thinks this line of thinking is important.

²²This point lies at the intersection of phonetics and phonology, but many good phonology textbooks will make it clear. See, for instance, Kenstowicz (1994).

do not intend specific aspects of their phonology, beyond the general intention to speak. We cannot intentionally control our ability to understand speech, which is automatic. Yet such intentions, I claim, are unnecessary. If each of us, as a speaker, processes speech sounds roughly the same way, communication succeeds. It is not necessary to have any communicative intentions surrounding that ability, beyond the general intention to speak and be understood. Our more general intentions to speak and be understood are fully served by having phonetic and phonological abilities, and others having overlapping ones. If they do, we succeed, if they do not, we may fail. I do not see any additional role for intentions here.

The same can be said for the more abstract aspects of syntax; especially, functional categories and the grammar they bring with them. As I already noted, we do not have anything like regular intentions going along with these. I doubt very much we have anything like communicative intentions corresponding to them, again except the very general one to 'say something'. As with sounds, overlap seems to me to suffice. If my hearer and I share enough grammar, communication succeeds and our communicative intentions are satisfied; if not, it fails. But no communicative intentions need to attach to functional aspects of grammar. Like the processing speech sounds, our understanding of grammar happens automatically, and either it overlaps across speakers or it does not.

The case I made above is that setting the standard relies on these kinds of abilities, and also other abilities like our ability to do tacit statistics, which are not fundamentally communicative at all. That is what I argued for at length. I now add the claim that these can affect communication by merely overlapping across speakers, and we need not, and should not, look for specific communicative intentions to which they are linked. I have not fully supported this claim, by any means. Even so, I hope to have made clear what the two general views of how lower-level cognitive abilities support communication are, and I have offered a few illustrative considerations about why one might prefer my non-intentional option.

I claimed that the less intentional indirect metasemantics I have offered is more accurate than an intention-based one such as King's. If we accept the view of communication I just endorsed, I think this is plausible, based on the available evidence. If we do not accept this view of communication, and insist on a more intentional view of communication, then the thoroughly intentional option seems required for the metasemantics as well. I thus make a conditional claim about the status of intentional indirect metasemantics. If you opt for a view where underlying cognitive abilities can serve communication merely by overlap, then my less intentional version of the metasemantics is clearly preferable. It offers more detail and a more refined account of the facts, and does not look for intentions where they are not to be found. On the other hand, if you insist that only shared intentions can serve communication, then King's intentional version is preferable. But either way, the metasemantics for parameters like the standard must be indirect.

4 Conclusion

I have argued in favor of indirect metasemantics for functional parameters like the standard. We can identify multiple influences on the standard, including comparison classes, lexical content, overt discourse, speaker intentions, and so on. These must be combined in context, and in different contexts, they can combine differently. This makes the case for indirectness. Furthermore, I argued, a uniformly intentional metasemantics like King's must also be indirect, if it is to handle a full range of cases. Thus, indirectness is not a specific feature of my preferred form of metasemantics.

I also considered the role of intentions in metasemantics. I did note that there is a substantial role for intentions, and conceded that this role may go beyond what I suggested in earlier work. But, I argued against the kind of thoroughly intentional option King offers. I argued this mainly on cognitive grounds. The inputs to our cognition of standard values is multi-faceted. I argued that it relies on inputs from grammatical competence (both syntactic and semantic), various ways we can recognize magnitudes in our environment, how we are sensitive to statistical regularities, and so on. It also relies on our communicative intentions, but many of the factors on this list are not best described as intentions.

Finally, I noted that which option is right, mine or King's intentional option, depends on some very large-scale questions about how communication works. I sketched a view of communication that allows many of the cognitive abilities I just mentioned to serve communication by mere overlap, and not via intentions at all. I did not fully justify that view, and I take it to be an open question which view is really right. But I do offer a conditional claim about how intentional a metasemantics should be. If you opt for my view of communication, you should also opt for my less intentional metasemantics.

In closing, I would like to take up, briefly, one further issue that is sometimes raised as an objection against an indirect metasemantics. Saying that somehow, in context, multiple factors get integrated, may seem like a weak theory, that fails to explain what is actually happening. I take this to indicate a set of open questions, rather than an as objection to an indirect metasemantics. If I am right about the range of factors that go into setting functional parameters—both linguistic and cognitive factors—then immediate or general answers about how they combine are not likely to be found. Many of the factors are already difficult to describe individually, and how they integrate in communicative settings is not an easy matter. We have now seen that big-picture questions about the nature of communication must be added to this list of open problems. The indirect metasemantics is an initial characterization of how functional parameters are set in context, and indicates some of the issues and problems that need to be addressed to fully describe this process. Most of these problems remain open.

References

- Abney, S. P. (1987). *The English noun phrase in its sentential aspect*. Ph.D. dissertation, Massachusetts Institute of Technology.
- Bach, K. (1994). Conversational implicature. *Mind and Language*, 9, 124–162.
- Baker, M. C. (2003). *Lexical categories*. Cambridge: Cambridge University Press.
- Barker, C. (2002). The dynamics of vagueness. *Linguistics and Philosophy*, 25, 1–36.
- Barner, D., & Snedeker, J. (2008). Compositionality and statistics in adjective acquisition: 4-year-olds interpret *tall* and *short* based on the size distributions of novel noun referents. *Child Development*, 79, 594–608.
- Bartsch, R., & Vennemann, T. (1973). *Semantic structures: A study in the relation between semantics and syntax* (2nd ed.). Frankfurt: Athenäum.
- Bierwisch, M. (1989). The semantics of gradation. In M. Bierwisch & E. Lang (Eds.), *Dimensional adjectives* (pp. 71–261). Berlin: Springer.
- Bogusławski, A. (1975). Measures are measures: In defense of the diversity of comparatives and positives. *Linguistische Berichte*, 36, 1–9.
- Burnett, H. (2014). A delineation solution to the puzzles of absolute adjectives. *Linguistics and Philosophy*, 37, 1–39.
- Cantlon, J. F., Platt, M. L., & Brannon, E. M. (2009). Beyond the number domain. *Trends in Cognitive Science*, 13, 83–91.
- Carey, S. (2009). *The origin of concepts*. Oxford: Oxford University Press.
- Cresswell, M. J. (1977). The semantics of degree. In B. Partee (Ed.), *Montague grammar* (pp. 261–292). New York: Academic.
- Cruse, D. A. (1986). *Lexical semantics*. Cambridge: Cambridge University Press.
- Dehaene, S. (2011). *The number sense: How the mind creates mathematics* (2nd ed.). Oxford: Oxford University Press.
- Fara, D. G. (2000). Shifting sands: An interest-relative theory of vagueness. *Philosophical Topics*, 20, 45–81. Originally published under the name ‘Delia Graff’.
- Feigenson, L. (2007). The equality of quantity. *Trends in Cognitive Science*, 11, 185–187.
- Fodor, J. A. (1983). *The modularity of mind*. Cambridge: MIT Press.
- Frank, M. C., Goodman, N. D., & Tenenbaum, J. B. (2009). Using speakers’ referential intentions to model early cross-situational world learning. *Psychological Science*, 20, 578–585.
- Frazier, L., Clifton Jr., C., & Stolterfoht, B. (2008). Scale structure: Processing minimum standard and maximum standard scalar adjectives. *Cognition*, 106, 299–324.
- Fukui, N. (2001). Phrase structure. In M. Baltin & C. Collins (Eds.), *Handbook of contemporary syntactic theory* (pp. 374–406). Oxford: Blackwell.
- Gallistel, C. R., & Gelman, R. (1992). Preverbal and verbal counting and computation. *Cognition*, 44, 43–74.
- Glanzberg, M. (2007). Context, content, and relativism. *Philosophical Studies*, 136, 1–29.
- Glanzberg, M. (2008). Metaphor and lexical semantics. *The Baltic International Yearbook of Cognition, Logic and Communication*, 3, 1–47. <https://newprairiepress.org/biyclc/vol3/iss1/2/>.
- Glanzberg, M. (2014). Explanation and partiality in semantic theory. In A. Burgess & B. Sherman (Eds.), *Metasemantics: New essays on the foundations of meaning* (pp. 259–292). Oxford: Oxford University Press.
- Glanzberg, M. (2016). Not all contextual parameters are alike. Unpublished manuscript.
- Grimshaw, J. (2005). Extended projection. In *Words and structure* (pp. 1–73). Stanford: CSLI Publications.
- Heim, I. (1985). Notes on comparatives and related matters. Unpublished manuscript, University of Texas, Austin. Available at <http://semanticsarchive.net/Archive/zc0ZjY0M/>.
- Heim, I. (1992). Presupposition projection and the semantics of attitude verbs. *Journal of Semantics*, 9, 183–221.
- Kaplan, D. (1989a). Afterthoughts. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan* (pp. 565–614). Oxford: Oxford University Press.

- Kaplan, D. (1989b). Demonstratives. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan* (pp. 481–563). Oxford: Oxford University Press. First publication of a widely circulated manuscript dated 1977.
- Kennedy, C. (1997). *Projecting the adjective*. Ph.D. dissertation, University of California, Santa Cruz. Published by Garland, 1999.
- Kennedy, C. (2007). Vagueness and grammar: The semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy*, 30, 1–45.
- Kennedy, C., & McNally, L. (2005). Scale structure, degree modification, and the semantics of gradable predicates. *Language*, 81, 345–381.
- Kenstowicz, M. (1994). *Phonology in generative grammar*. Cambridge: Blackwell.
- King, J. C. (2013). Supplementives, the coordination account, and conflicting intentions. *Philosophical Perspectives*, 27, 288–311.
- King, J. C. (2014a). The metasemantics of contextual sensitivity. In A. Burgess & B. Sherman (Eds.), *Metasemantics: New essays on the foundations of meaning* (pp. 97–118). Oxford: Oxford University Press.
- King, J. C. (2014b). Speaker intentions in context. *Nous*, 48, 219–237.
- Klein, E. (1980). A semantics for positive and comparative adjectives. *Linguistics and Philosophy*, 4, 1–45.
- McMurray, B., Horst, J. S., & Samuelson, L. K. (2012). Word learning emerges from the interaction of online referent selection and slow associative learning. *Psychological Review*, 119, 831–877.
- Meck, W. H., & Church, R. M. (1983). A mode control model for counting and timing processes. *Journal of Experimental Psychology: Animal Behavior Processes*, 9, 320–334.
- Morzycki, M. (2016). *Modification*. Cambridge: Cambridge University Press.
- Perry, J. (1996). Thought without representation. *Proceedings of the Aristotelian Society Supplementary Volume*, 60, 263–283.
- Recanati, F. (2004). *Literal meaning*. Cambridge: Cambridge University Press.
- Rett, J. (2015). *The semantics of evaluativity*. Oxford: Oxford University Press.
- Rips, L. J., & Turnbull, W. (1980). How big is big? Relative and absolute properties in memory. *Cognition*, 8, 145–174.
- Roberts, C. (1996). Information structure in discourse: Towards an integrated formal theory of pragmatics. *Ohio State University Working Papers in Linguistics*, 49, 91–136.
- Schmidt, L. A., Goodman, N. D., Barner, D., & Tenenbaum, J. B. (2009). How tall is tall? Compositionality, statistics, and gradable adjectives. *Proceedings of the Annual Conference of the Cognitive Science Society*, 31, 3151–3156.
- Smith, L. B., Suanda, S. H., & Yu, C. (2014). The unrealized promise of infant statistical word-referent learning. *Trends in Cognitive Science*, 18, 251–258.
- Solt, S., & Gotzner, N. (2012). Experimenting with degrees. *Proceedings of Semantics and Linguistic Theory*, 22, 166–187.
- Sperber, D., & Wilson, D. (1998). *Relevance*. Cambridge: Harvard University Press.
- Stalnaker, R. C. (1978). Assertion. In P. Cole (Ed.), *Pragmatics* (Syntax and semantics, Vol. 9, pp. 315–322). New York: Academic. Reprinted in Stalnaker (1999).
- Stalnaker, R. C. (1997). Reference and necessity. In B. Hale & C. Wright (Eds.), *Companion to the philosophy of language* (pp. 534–554). Oxford: Blackwell.
- Stalnaker, R. C. (1999). *Context and content*. Oxford: Oxford University Press.
- Stanley, J. (2000). Context and logical form. *Linguistics and Philosophy*, 23, 391–434.
- Stanley, J., & Szabó, Z. G. (2000). On quantifier domain restriction. *Mind and Language*, 15, 219–261.
- Stevens, S. S. (1975). *Psychophysics: Introduction to its perceptual, neural, and social prospects*. New York: Wiley. Edited by G. Stevens.
- Syrett, K., Kennedy, C., & Lidz, J. (2010). Meaning and context in children’s understanding of gradable adjectives. *Journal of Semantics*, 27, 1–35.
- von Stechow, A. (1984). Comparing semantic theories of comparison. *Journal of Semantics*, 3, 1–77.

- Winter, Y., & Rotstein, C. (2004). Total adjectives vs. partial adjectives: Scale structure and higher-order modification. *Natural Language Semantics*, *12*, 259–288.
- Xu, F., & Tenenbaum, J. B. (2007). Word learning as Bayesian inference. *Psychological Review*, *114*, 245–272.
- Yu, C., & Smith, L. B. (2012). Modeling cross-situational word-referent learning: Prior questions. *Psychological Review*, *119*, 21–39.

Speaker Intentions and Objective Metasemantics



Jeffrey C. King

Abstract The word ‘I’ has a context invariant meaning that suffices to determine semantic values for it in context. Many people, including the author of the present paper, think that a lot of contextually sensitive expressions aren’t like that: their context invariant meanings *don’t* by themselves suffice to secure semantic values for them in context. Demonstratives and demonstratively used pronouns are familiar examples of expressions that appear to be of this sort. But arguably, tense, quantifiers, gradable adjectives, modals, conditionals, possessives, expressions that take implicit arguments (‘John has had enough’), ‘only’, and more are examples of expressions of this sort as well. For such expressions, the question arises as to the mechanism by means of which they secure semantic values in context. Call an account of how one of these expressions manages to secure a semantic value in context a *metasemantics* for the expression in question.

As I suggested above, I believe that the context invariant meanings of the contextually sensitive expressions just mentioned must be supplemented in some way in context in order for these expressions to have semantic values in context. For this reason, I call these expressions *supplementives*. In a series of recent papers, I have defended a single metasemantics that I claim applies to all supplementives. I call my metasemantics for supplementives the *coordination account*.

In a recent paper Michael Glanzberg argues against the coordination account and in favor of a very different metasemantics. The purpose of the present paper is to respond to Glanzberg’s arguments against the coordination account and in favor of his metasemantics.

Keywords Speaker intentions · Contextual sensitivity · Metasemantics · Gradable adjectives · Demonstratives

J. C. King (✉)

Department of Philosophy, Rutgers University, New Brunswick, NJ, USA

e-mail: jking@philosophy.rutgers.edu

© Springer Nature Switzerland AG 2020

T. Ciecierski, P. Grabarczyk (eds.), *The Architecture of Context and Context-Sensitivity*, Studies in Linguistics and Philosophy 103, https://doi.org/10.1007/978-3-030-34485-6_4

The word ‘I’ has a context invariant meaning that suffices to determine semantic values for it in context. Many people, including the author of the present paper, think that a lot of contextually sensitive expressions aren’t like that: their context invariant meanings *don’t* by themselves suffice to secure semantic values for them in context. Demonstratives and demonstratively used pronouns are familiar examples of expressions that appear to be of this sort. But arguably, tense, quantifiers, gradable adjectives, modals, conditionals, possessives, expressions that take implicit arguments (‘John has had enough’), ‘only’, and more are examples of expressions of this sort as well. For such expressions, the question arises as to the mechanism by means of which they secure semantic values in context. Call an account of how one of these expressions manages to secure a semantic value in context a *metasemantics* for the expression in question.

As I suggested above, I believe that the context invariant meanings of the contextually sensitive expressions just mentioned must be supplemented in some way in context in order for these expressions to have semantic values in context. For this reason, I call these expressions *supplementives*. In a series of recent papers, I have defended a single metasemantics that I claim applies to all supplementives.¹ I call my metasemantics for supplementives the *coordination account*:

Coordination Account Metasemantics

A speaker S’s use δ of a supplementive in context c has o as its semantic value iff 1. S intends o to be the semantic value of δ in c ; and 2. a competent, reasonable, attentive hearer H who knows the common ground of the conversation at the time S utters δ , and who has the properties attributed to the audience by the common ground at the time S utters δ would know² that S intends o to be the semantic value of δ in c .³

I’ll sometimes abbreviate condition 2 by saying that *an idealized hearer who knows the common ground of the conversation would know that S intends o to be the semantic value of δ in c* . When both conditions in the coordination account are met, say that *S has a recognizable intention that o be the semantic value in c of δ* . So according to the coordination account, recognizable intentions fix the semantic values in context of supplementives.

In a recent paper Michael Glanzberg argues against the coordination account and in favor of a metasemantics on which numerous features of a conversation determine the semantic values of supplementives in context.⁴ According to Glanzberg, shared presuppositions, comparison classes determined by salient groups, overt prior

¹King (2013, 2014a, b).

²Perhaps what is required here is that the hearer would bear an attitude towards the claim that S intends o to be the semantic value of δ in c that is just like knowledge except for lacking safety (assuming safety *is* required for knowledge). See King (2013) note 5 for discussion.

³This is the version of the coordination account I call *Bad Intentions* in King (2013). The requirement in condition 2 that the hearer have the properties attributed by the common ground to the audience will play almost no role here. See King (2013) for discussion. I actually now favor the version of the coordination account I there call *Best Laid Plans*, but the latter is more complex and the additional complexity won’t be relevant to the examples in the present paper.

⁴Glanzberg (2019). All references are to this paper unless otherwise indicated.

discourse, speaker intentions and other features determine the semantic values in context of supplementives. Because Glanzberg's metasemantics claims that things other than speaker intentions play a role in securing semantic values for supplementives in context and because many of them are objective features of the conversation, I call metasemantic accounts like Glanzberg's *objective* metasemantics. The purpose of the present paper is to respond to Glanzberg's arguments against the coordination account and in favor of his objective metasemantics.

The only supplementives Glanzberg considers are gradable adjectives. Following Glanzberg, I'll assume a Kennedy (2007) style semantics for gradable adjectives on which they denote *measure functions*: functions that map individuals to degrees (type $\langle e, d \rangle$). In turn, these degrees are totally ordered with respect to some dimension given by the adjectives meaning (e.g. height in the case of 'tall'), yielding a *scale*.⁵ Adjectives combine with degree morphology yielding expressions that denote properties of or relations between individuals. In the positive form ('is tall'), the degree morpheme is a null morpheme *pos*. Hence, syntactically, the positive form of the adjective with degree morpheme looks as follows:

1. [DegP[[Deg*pos*] [Aptall]]]

The semantics for *pos* is: $\|pos\|^c = \lambda g \lambda x. g(x) > d_c$, where d_c is a degree of tallness determined in context c and g ranges over adjective meanings (measure functions). Thus $\|[[[DegP[[Deg*pos*] [Aptall]]]]\|^c = \lambda x. \mathbf{tall}(x) > d_c$, where **tall** is the semantic value of 'tall'.⁶ An individual o has this property just in case the height **tall** assigns to o is greater than the degree of height d_c determined in the context c . So a sentence like

2. LeSean is tall.

is true in context c just in case

2TC. **tall** (LeSean) $> d_c$

On a theory like this, the gradable adjective, or really, *pos*, is the supplementive and a degree on the relevant scale is the semantic value it gets assigned in context (d_c in the present case). The coordination account claims it gets assigned a degree d_c in context c by the speaker recognizably intending d_c to be the semantic value in c of *pos*.⁷

⁵Kennedy (2007) doesn't commit to any particular way of formalizing scales, but he says that they "minimally" have to be triples $\langle D, <, \delta \rangle$, where D is a set of points, $<$ totally orders D and δ is a dimension (wealth, height, etc.). I assume Kennedy means that $<$ *strictly totally orders* D , that is, that $<$ is irreflexive, asymmetric, transitive and connected.

⁶Here I depart slightly from Kennedy's (2007) view, on which $\|pos\|^c = \lambda g \lambda x. g(x) > s(\mathbf{tall})$, where s is a contextually determined function from adjective meanings to degrees. For various reasons, I prefer the view in the text.

⁷Of course ordinary speakers wouldn't describe their intentions this way since they aren't aware of *pos*. This is how we as theorists describe the relevant intentions. Speakers *would* describe their intentions as determining a cutoff in context for e.g. 'tall'. For some complications regarding this, see King (2014a). Further, in some felicitous uses of gradable adjectives in the positive form,

Earlier I said that Glanzberg argues against the coordination account, but that isn't quite right. Glanzberg calls supplementives *contextual parameters* and he distinguishes two types: *thematic parameters* and *functional parameters*.⁸ Glanzberg is happy to accept the coordination account for demonstratives, demonstrative pronouns and thematic parameters generally, or at least he is happy not to challenge it for these supplementives.⁹ However, for covert functional parameters, covert parameters occurring syntactically in functional categories, Glanzberg rejects the coordination account and argues for his objective metasemantics. Recall from above that *pos* is the supplementive/contextual parameter for gradable adjectives in the positive form. Recall also that it occurs in *Deg/Deg_p*, which are functional categories. Hence, *pos* is a covert functional parameter and as such Glanzberg thinks it requires his objective metasemantics. My view, of course, is that the coordination account is the correct metasemantics for *pos*. Let's turn to Glanzberg's arguments that *pos* requires an objective metasemantics instead of the coordination account. The dialectic here is a bit convoluted, because in the end the arguments for his objective metasemantics we are about to consider aren't the main arguments for his view. He later seems to admit that the coordination account has responses to his arguments here and gives a different argument against the coordination account in the end. However, Glanzberg seems to at least think that his arguments here make a *prima facie* case for his objective metasemantics, so it is worth considering them and how an advocate of the coordination account can best respond to them. To this task we now turn.

As I indicated above, Glanzberg thinks that a number of factors, together or separately, secure semantic values for supplementives in context. He calls a metasemantics like this, where several different factors, alone or together, secure semantic values for supplementives in context an *indirect metasemantics*. *Prima facie*, an account like the coordination account on which a single factor—recognizable speaker intentions—secures semantic values for supplementives in context is what Glanzberg calls a *direct metasemantics*. However, as we'll see later Glanzberg argues that the coordination account is itself an indirect metasemantics. The arguments we are now going to consider in favor of Glanzberg's objective metasemantics consist of a series of cases in which Glanzberg claims that various different factors secure semantic values for supplementives in context in his different cases. As noted above, all Glanzberg's cases involve gradable adjectives. So in each case, *pos* gets assigned a degree in context. Glanzberg talks about these cases in terms of *setting* or *determining the standard* (a degree: *d_c*) *in context*. In his terminology, then, the cases are supposed to show that different factors set the standard (in context) in different cases.

speakers don't have intentions that determine *unique* cutoffs/values for *pos* in context. See King (2018) for discussion.

⁸See Glanzberg (2007, 2016, 2019).

⁹P. 15.

Here is Glanzberg's first case. He had been discussing comparison classes setting the standard, which we will return to below, and here he claims *lexical meaning* sets the standard:

Following Kennedy and McNally (2005), we can observe another important factor in some cases. The scale structure that an adjective lexically encodes can be crucial. Kennedy and McNally observe that 'absolute' adjectives like 'open' and 'straight', which have a minimum or maximum value for their scale, at least usually require the extreme value to be taken. Consider:

- (10) a. The door is open.
b. The rod is straight.

To be straight is to be completely straight, and so to have the maximum value on the scale. To be open is to be even a small amount open, and so to have above the minimum value on the scale. Crucially, we see little or no context dependence in absolute cases. The reason is that the lexical meaning of the adjective does all the work, and nothing else about the context (typically) matters. Hence, comparison classes can be overruled by lexical meaning.

We thus have at least two factors that contribute to setting the standard: comparison class, and lexical meaning. Both can contribute, but neither is sufficient to fix standards for all adjectives in all contexts.¹⁰

Glanzberg follows Kennedy and McNally in claiming that absolute adjectives in virtue of their meanings take the maximum or minimum degree on the relevant scales as their semantic values in context. 'Straight' takes the maximum degree of straightness on its associated scale (to be straight is to be completely straight) and 'open' takes the minimum degree of openness on its associated scale. But as Glanzberg himself seems to say, that means we really don't have context sensitivity here at all. Or if we do, we should think of this as a case like the pure indexical 'I' where the context invariant meaning of the contextually sensitive expression suffices by itself to secure a semantic value for the expression in context (an absolute gradable adjective on this view would have the unusual feature of being a contextually sensitive expression that has the same semantic value in every context). But either way, such absolute gradable adjectives are not supplementives at all. That means that a metasemantics for supplementives is not relevant to them. As a result, this case is simply irrelevant to the debate over the correct metasemantics for supplementives. Hence, it provides no support for Glanzberg's objective metasemantics.

Here is Glanzberg's second case, in which he claims *shared presuppositions* set the standard/assign a semantic value in context to *pos*:

We thus have at least two factors that contribute to setting the standard: comparison class, and lexical meaning. Both can contribute, but neither is sufficient to fix standards for all adjectives in all contexts.

I have argued that other factors can play a role too. Most of the kinds of things that influence context can. Not surprisingly, shared presuppositions can (as we learned from Stalnaker (1978) among many others). Consider:

¹⁰Pps. 6–7.

(11) Context: We are at a summit on international development. Shared presuppositions about the problems facing poor nations, and the things that can be done to fix them, make it clear that individuals or countries can only make an impact with contributions of \$10 million or more. I make manifest my intention to count average philosophers as rich, by holding up photos of conferences and saying ‘look at all those rich people’.

Assertion: We should try to get money from philosophers because they are rich.

Absent any other clear contextual information, the claim that philosophers are rich just appears false in this context (though of course, it could be true in a context where we were talking about, say, relative income in the United States). In this case, the shared presuppositions about our purposes in classifying as rich beat out an attempt to make a salient comparison class, and they beat out a fairly clear communicative intention.¹¹

I agree with Glanzberg that the claim that philosophers are rich does not seem true here. I also think the entire assertion itself does not seem true. This is what the coordination account and Glanzberg’s objective metasemantics must explain.

Glanzberg claims that the shared presuppositions about our purposes in classifying people as rich in this conversation set the semantic value in context of *pos* very high on the wealth scale in such a way that the claim that philosophers are rich comes out false. He doesn’t attempt to explain the fact that the entire assertion seems untrue or false but that would seem to follow from the fact that the claim that philosophers are rich is false: the truth of sentences of the form ‘A because B’ presumably requires B’s truth.

Glanzberg seems to think that an account like the coordination account on which recognizable speaker intentions fix the semantic value in context of *pos* will incorrectly predict that ‘they are rich’ expresses something true here. However, I don’t think that’s right. For the conditions of the coordination account are not met here. For those conditions to be met, (1) the speaker S must intend a degree *d* of wealth as the threshold for being rich; and (2) a competent, reasonable, attentive hearer H who knows the common ground of the conversation at the time S utters ‘rich’, and who has the properties attributed to the audience by the common ground at the time S utters ‘rich’ would know that S intends *d* to be the threshold for being rich in *c*. Though Glanzberg talks about a “fairly clear communicative intention” here, it isn’t at all clear that *either* condition is met in this example. Glanzberg doesn’t mention that the speaker intends a specific degree on the wealth scale to be the semantic value of *pos* in context here. Let’s suppose the speaker intended having an annual income of \$100,000 to be the threshold for being rich so that condition 1 of the coordination account is satisfied. Call this *version 1 of Glanzberg’s ‘rich’ example*. Still, holding up pictures of people at conferences and saying ‘look at these rich people’ won’t allow a competent, reasonable, attentive hearer H who knows the common ground of the conversation at the time S utters ‘rich’, and who has the properties attributed to the audience by the common ground at the time S utters ‘rich’ to know that S intends an annual income of \$100,000 to be the threshold for being rich in *c*. How would such a hearer know how much money the people in the pictures make? So condition 2 of the coordination account is not met with

¹¹p. 7.

the result that far from predicting ‘they are rich’ expresses something true here, the coordination account predicts that it expresses something not truth evaluable, and hence not true. So the coordination account predicts that the claim that philosophers are rich is not true here.

The coordination account also predicts that the entire assertion is not true:

3. We should try to get money from philosophers because they are rich.

3 is of the form ‘A because B’ and as indicated above, for something of this form to be true, B must be true. But we just saw that the coordination account predicts that B expresses something not true. So then it predicts 3 as a whole is not true. So the coordination account claims that ‘they are rich’ expresses something not true as the example is constructed, whereas Glanzberg can claim it expresses something false. Both accounts predict that 3 is either false or not true. I don’t see that Glanzberg’s explanation of the case as it stands has any advantage over the coordination account’s explanation. In particular, both accounts can claim ‘they are rich’ isn’t true. I don’t think being able to say it is false is any real advantage, mainly because I am skeptical that people clearly distinguish between not being true and being false in making judgments about such cases.

Perhaps we can better compare Glanzberg’s account and the coordination account by altering the case as little as we can while insuring both conditions of the coordination account are satisfied. Call this *version 2 of Glanzberg’s ‘rich’ example*. Imagine that the speaker has pictures of people where under each picture it says: ‘Philosophy Professor. Annual income: \$101,000’ (or some other number over \$100,000). The speaker then holds up the pictures and says: ‘Look at these rich people. They all make over \$100, 000 per year’. This would appear to satisfy condition 2 of the coordination account, so that now both conditions are satisfied and \$100,000 is the threshold for being rich in this context according to the coordination account. Suppose that 3 is uttered next.

In this case 3 seems false as before and the claim that philosophers are rich sounds infelicitous or inappropriate. I don’t think it sounds clearly untrue in this case, as it is clear to all that the speaker is conveying the true claim that philosophers are rich_{100,000}.¹² So what needs to be explained here is why 3 sounds false and why the claim that philosophers are rich sounds infelicitous or inappropriate.

If we assume that ‘rich’ here gets assigned the threshold of having an annual income of \$100,000 as the coordination account claims, ‘they are rich’ should sound infelicitous in 3. After all, that philosophers are rich_{100,000} in 3 is being given as a reason for thinking we should get money from them to fix the problems facing poor nations, where only someone who can give \$10 million or more can make an

¹²Following Karen Lewis in unpublished work, I use ‘Philosophers are rich_{100,000}’ to express the proposition expressed by ‘Philosophers are rich’, where the threshold for ‘rich’ is fixed at \$100,000. All parties should agree that in this case the speaker is *conveying* the claim that philosophers are rich_{100,000} in uttering ‘they are rich’. The only question is whether the speaker *semantically expresses* this claim in speaking as he does in this context.

impact. But clearly, that philosophers are rich_{100,000} is a *terrible* reason for thinking we should get money from them to fix the problems facing poor nations, where only someone who can give \$10 million or more can make an impact.¹³ This also means that 3 is false if ‘they are rich’ expresses the true claim that philosophers are rich_{100,000} as the coordination account predicts, because it thereby is a terrible reason for thinking we should get money from philosophers to help poor nations needing at least \$10 million and yet it is being given as a reason for such.¹⁴

So version 2 of Glanzberg’s ‘rich’ example can also be explained by the coordination account even given that it claims ‘they are rich’ in 3 expresses the true claim that philosophers are rich_{100,000}. For it predicts that ‘they are rich’ as it occurs in 3 is infelicitous. Of course Glanzberg predicts ‘they are rich’ to be false as it occurs in 3. But, again, I have a hard time seeing how that is a significant advantage for Glanzberg. As long as both accounts predict there is something “bad” about ‘they are rich’ as it occurs in 3, both accounts capture the judgment that something went wrong with ‘they are rich’ as it occurs in 3. In a moment, we’ll see another reason for thinking ‘they are rich’ should seem infelicitous in this case if the coordination account is right that it expresses the claim that philosophers are rich_{100,000}. Further, both accounts predict that 3 is false. Hence version 2 of Glanzberg’s ‘rich’ example as it stands provides no argument against the coordination account and no argument for Glanzberg’s objective metasemantics and the claim that the standard is set in this case by shared presuppositions. I should add that I think the falsity of 3 produces noise in judging whether ‘they are rich’ is true or false in 3 and further disinclines speakers to judge it to be true in 3.

So let’s change the example so that ‘they are rich’ is not being given as a reason for thinking we should get money from philosophers to help fix the problems of poor nations to eliminate that bit of noise that I claim disinclines speakers to judge that ‘they are rich’ in 3 is true. Call this *version 3 of Glanzberg’s ‘rich’ example*. Hence, suppose after holding up the photos of philosophers, where under the photo it states their profession and salary, and saying ‘These people all make over \$100,000 per year’, the speaker simply says

4. Philosophers are rich.

Here both conditions of the coordination account are met and it predicts that 4 expresses the proposition that philosophers are rich_{100,000}. Clearly, there still is something bad or wrong about uttering 4 in this situation. But if the coordination account is right, 4 expresses the truth that philosophers are rich_{100,000}. Further, it is not being given as a bad reason for thinking we should try to get money from philosophers to help poor countries. So why does 4 sound bad or wrong *here* if

¹³If I say to you out of the blue ‘I am moving to New York, because I love the west coast’, ‘I love the west coast’ seems odd or infelicitous even if true because it provides such a bad reason for moving to New York absent other considerations.

¹⁴I assume here that a sentence of the form ‘A because B’ is false if B is true but is a terrible reason for A.

the coordination account is right? The answer is not far to seek. Given that we are in a conversation in which we are interested in individuals who can afford to donate \$10 million dollars, setting the threshold for being rich at having an annual income of \$100,000 is inappropriate. It renders the utterance of 4 irrelevant, and inappropriate for current concerns, and so infelicitous. After all, why is the speaker discussing who is rich_{100,000} in a context where we need people capable of donating \$10 million? So even if 4 expresses the true claim that philosophers are rich_{100,000} as the coordination account claims in version 3 of Glanzberg's example, we can explain why it sounds bad and infelicitous. Glanzberg can claim instead that shared presuppositions set the standard very high on the wealth scale so that 4 is false, but it isn't at all clear whose explanation of the badness or wrongness of 4 wins out in the end. So here again I think we have a stalemate between the coordination account and Glanzberg's objective semantics. But obviously that does no good for the defense of Glanzberg's account. It is worth adding that the coordination account's explanation of why 4 is infelicitous here will also provide another reason for why 'they are rich' is infelicitous in 3 in version 2 of Glanzberg's 'rich' example. For there too, in setting the standard at \$100,000 the speaker is using 'they are rich' to express something irrelevant and inappropriate given conversational purposes, and so the sentence is infelicitous.

Finally, it is worth noting that if we change Glanzberg's example minimally, so that setting the standard at \$100,000 *isn't* inappropriate and irrelevant, 4 suddenly becomes felicitous and seems true. Call this *version 4 of Glanzberg's 'rich' example*. Imagine that the summit on international development is winding down and we are getting ready to go back home. Holding up the pictures of philosophers as before, where under each photo the profession and salary of the individual pictured is stated, I say:

4+. We have been talking about the ultrawealthy during this summit. As we go back to our communities to work with the poor there, we shouldn't focus on the ultrawealthy. All of these people make over \$100,000 dollars. That's a lot of money. So as you go back to your communities, it is important to remember that philosophers are rich and could do a lot to help the poor.

Here the claim that philosophers are rich_{100,000} seems both felicitous and true, just as the coordination account predicts, given the change in conversational purposes. Of course Glanzberg could appeal to the overt discourse here as an objective factor that changes how the standard is set.

Because the dialectic has been complicated, let me summarize what the coordination account and Glanzberg's objective metasemantics claim about the different versions of Glanzberg's 'rich' example. In version 1, Glanzberg claims 'they are rich' is false. The coordination account claims it is not true. Both Glanzberg and the coordination account claim that 3 is false or untrue. Given the subtle difference between being not true and being false, I think neither theory has a clear advantage here. In version 2, the coordination account claims 'they are rich' in 3 is true but infelicitous for two reasons. First, it is being given as a reason for thinking we should try to get money from philosophers and that philosophers are rich_{100,000} is a terrible

reason for thinking we should try to get money from them, given that only people who can donate \$10 million can have any impact. Second, given that our concern is with individuals who can donate \$10 million, setting the threshold for being rich at \$100,000 is inappropriate and irrelevant. This makes ‘they are rich’ in 3, which according to the coordination account expresses the claim that philosophers are rich_{100,000}, itself inappropriate, irrelevant and so infelicitous. Here, the coordination account predicts 3 is false, as it seems to be. I claim this further disinclines speakers to judge that ‘they are rich’ in it expresses something true. Again, Glanzberg predicts that ‘they are rich’ expresses something false here due to shared presuppositions setting the standard very high. Again, I see no advantage for either account here. In version 3, the coordination account predicts 4 is true but infelicitous. Given that we are in a conversation in which we are interested in individuals who can afford to donate \$10 million dollars, setting the threshold for being rich at having an annual income of \$100,000 is inappropriate. It renders the utterance of 4 irrelevant, inappropriate and so infelicitous. Here again, Glanzberg predicts that 4 is false. As before, this seems to me a stalemate. Both accounts predict 4 will sound bad and wrong. Finally, version 4 of Glanzberg’s ‘rich’ example bolsters the explanation that the coordination account gave of version 3. It was claimed that in version 3, given conversational purposes and that 4 expressed the claim that philosophers are rich_{100,000} as the coordination account claims, 4 was inappropriate, irrelevant and so infelicitous. When we minimally changed conversational purposes in version 4, so that 4 expressing the claim that philosophers are rich_{100,000} is not inappropriate and irrelevant, 4 suddenly seemed felicitous and true. In the end, then, I don’t see that any version of Glanzberg’s ‘rich’ example provides support for his objective metasemantics over the coordination account.

Glanzberg’s next example purports to show that a *comparison class* can fix the standard:

In other cases, a comparison class can exert a stronger influence than common ground information that is presupposed by all parties. For instance:

(12) Context: We surrounded by basketball players, and everyone is talking about basketball. I happen to be obsessed with jockeys, and you know that I am. A 5’ 5” high man walks into the room. I say:

Assertion: He is tall.

It appears very hard to hear this as true. In this case, the standard is influenced by the salient group in the immediate environment, which can provide a comparison class. Our shared knowledge of my obsession with jockeys, and the communicative intention that goes with it, does not overrule this.¹⁵

I agree with Glanzberg’s judgment that it is hard to hear the assertion as true here. His view claims it is false.

For definiteness, suppose, as I think Glanzberg does, that the conversational participants are the speaker and the addressee (after all, it is important that the jockey obsession is common ground and only the addressee was claimed to know about it). So here Glanzberg claims that a salient group in the environment of

¹⁵Pps. 7–8.

the conversation—a group of basketball players—perhaps along with the subject matter of the conversation—basketball—provide a comparison class—the class of basketball players—that in turn sets the standard. He claims that this overrides the speaker’s intention. But note that again the coordination account conditions are not satisfied here. Even condition 1 appears not to be satisfied, since the speaker does not appear to intend a particular degree of height to be the semantic value of *pos* in context. Further, the fact that the speaker S is obsessed with jockeys does not guarantee that condition 2 is satisfied. That is, S’s known obsession with jockeys does not guarantee that a competent, reasonable, attentive hearer H who knows the common ground of the conversation at the time S utters ‘tall’, and who has the properties attributed to the audience by the common ground at the time S utters ‘tall’ would know that S intends the semantic value of *pos* in c to be the threshold for tall jockeys. I suppose if S is known to *always* use the threshold of being tall for a jockey whenever he uses ‘tall’ the condition would be satisfied. But in that case, it again is not obvious that ‘He is tall’ is false as opposed to merely irrelevant, inappropriate and so infelicitous given the context (“everyone is talking about basketball players”). This would explain Glanzberg’s judgment that “it is hard to hear this as true.” And if we give up the supposition that S always uses jockey height for his standard in using ‘tall’ and satisfy condition 2 some other way in the context of utterance, the judgment that ‘He is tall’ is false loses its force. Suppose, for example, as before S is obsessed with jockeys and their heights and this is known to the hearer. Further, S is clearly reading a picture-filled magazine called ‘People Who are 5’ Tall and Taller and so Would be Tall for a Jockey but not for a Basketball Player Quarterly’.¹⁶ Finally, suppose that though the basketball players are still salient to S, S clearly has no interest in them or the conversation he is hearing about them. Let all this be common ground. Now a 5’ 5” man walks in and S says ‘He’s tall.’ Here I think condition 2 of the coordination account is satisfied. But here, I find we can readily hear the utterance as true.¹⁷ So once we modify Glanzberg’s case here in such a way that the conditions of the coordination account are satisfied, the relevant utterance seems true just as the coordination account predicts. Further, in Glanzberg’s original version of the case, the coordination account predicts that no truth evaluable proposition was expressed. So we can actually agree with Glanzberg that in that version of the case it “appears very hard to hear this as true”. Hence, the case again provides no evidence for Glanzberg’s claim that a comparison class fixes the standard or for his objective metasemantics.

Glanzberg’s final case is intended to show that *overt prior discourse* can fix the semantic value of *pos* in context:

¹⁶I put the title this way, because ‘People who are tall for a jockey and not a basketball player’ seems to presuppose that they are jockeys and basketball players and we don’t want to assume that the man who is about to walk in is a jockey. Thanks for suggestions on the title to Annie Papreck King.

¹⁷Informants I have tried the example on agree.

Work in pragmatics has also shown how previous discourse can affect context in different ways than mere presupposition. We see that at work with standards too. If we supplement (12) with some discourse, the situation changes:

(13) Context: We surrounded [sic] by basketball players, and everyone is talking about basketball. I happen to be obsessed with jockeys, and you know that I am. A 5' 5" high man walks into the room.

Discourse: I am going to a horse race later today. I just know that the shorter jockey will win. They always do. You have to be under 5' to do really well.

Assertion: He is too tall.

I am not sure how firm the judgement is here, but it is much more easy to hear as true. So, overt discourse can override prior presuppositions, and can implicitly guide us to different comparison classes we well.¹⁸

So here Glanzberg claims that the prior discourse provides a comparison class—jockeys—that in turn secures a semantic value for *pos* (5') that makes 'He is too tall' true when 'He' picks out the 5'5" man that just walked in. By contrast, the coordination account claims that the prior discourse makes manifest the speaker's intention to make 5' the threshold for being tall in the context of utterance. As a result, both conditions of the coordination account are met. Clearly, the speaker intends 5' to be the threshold for being tall in this context. Further, precisely because of the previous discourse, a competent, reasonable, attentive hearer who knows the common ground of the conversation at the time the speaker utters 'tall', and who has the properties attributed to the audience by the common ground at the time S utters 'tall' would know that the speaker intends 5' to be the semantic value of *pos* in *c*. So the coordination account can explain why 5' is the semantic value in context of *pos* in this case and hence why the utterance is judged true. So Glanzberg's case fails to show that prior discourse fixes the semantic value of *pos* in context and so here again the case provides no argument in favor of Glanzberg's objective metasemantics.

This last response on the part of the coordination account brings up an important point. The coordination account is clearly not an objective metasemantics since it claims that semantic values of supplementives in context are fixed exclusively by recognizable speakers' intentions. However, there is a sense in which the coordination account mimics an objective metasemantics and allows that prior discourse, common ground, salient comparison classes and the like play an *indirect* role in fixing the semantic values of supplementives in context. We just saw an illustration of the way this occurs. Given the prior discourse, a speaker who wishes to be understood had better intend 5' to be the threshold for being tall in context. Further, given the prior discourse, a competent, reasonable, attentive hearer who knows the common ground of the conversation at the time the speaker utters 'tall', and who has the properties attributed to the audience by the common ground at the time the speaker utters 'tall' would know that the speaker intends 5' to be the semantic value of *pos* in *c*. In general many, many factors about the context of utterance, including the meanings of the words the speaker is uttering, prior discourse, questions under discussion, and all those mentioned by Glanzberg as

¹⁸P. 8.

factors in his objective metasemantics, *constrain* what a speaker can reasonably intend to be the semantic value of a use of a supplementive and *determine* that an idealized hearer who knows the common ground of the conversation would know what the speaker intends. Let me give two more examples to illustrate this point, because it will be important later. Suppose two of us are in the locker room of a professional basketball team and the only people present are us and the players. If I use ‘tall’ in this context and wish to be understood, I had better intend the threshold for ‘tall’ to be appropriate for basketball players unless I do something else to make a different intention recognizable. Similarly, in such a case absent any indication to the contrary, an idealized hearer will know that I intend such a threshold. So here again, an objective feature of the conversation—a salient group determining a comparison class—can constrain what semantic value a speaker can reasonably intend and determine what an idealized hearer who knows the common ground of the conversation would know about the speaker’s intention. For a final example, suppose that it is common ground that we are meeting Glenn for lunch at Gjusta, that we can all recognize his car and that as we pull up to Gjusta Glenn’s car is parked out front (we all see that we all see this). Pointing at the car, I say:

5. Looks like he beat us here.

In such a case, unless I make a different intention manifest in some way, if I wish to be understood I had better intend Glenn as the semantic value of ‘he’ in context and an idealized hearer who knows the common ground of the conversation would know that I so intend. So once again, we have objective factors—the common ground—constraining what a speaker who wishes to be understood can intend and determining what an idealized hearer who knows the common ground of the conversation would know about what the speaker intends.

One more point should be made before turning to other issues. Glanzberg often says things like the following:

A contextual parameter with an indirect metasemantics must be set by the various pieces of information context provides, but context does not simply hand us a value for such a parameter, nor does it hand us a uniform rule for computing the value from a specific piece of contextual information. Rather, a range of contextual information and computational rules must be taken into account and weighed in working out the value from context.¹⁹

This makes it sound as though he claims that various contextual factors—salient groups, comparison classes, shared presuppositions, overt discourse and so on—combine in different ways in different cases to fix the semantic value of *pos* in context. However, as we have seen when he gives examples that purport to show that these factors fix the semantic value of *pos* in context, only one factor fixes the semantic value in context. In his first example, it was alleged to be lexical meaning, in the second shared presuppositions, in the third, a salient group and in the fourth overt discourse. So Glanzberg hasn’t really offered a case that purports to show

¹⁹P. 9.

that various factors *together* get weighed and combined to yield a semantic value in context for *pos*.

In any case, Glanzberg argues next that my metasemantics is *indirect* in his sense. Recall that this means for Glanzberg that multiple factors play a role in fixing a semantic value for *pos* in context. This could mean that different factors play a role in different cases or it could mean that different factors may play a role in a single case. As we saw, though Glanzberg appears to claim the latter, the actual examples he gives only purport to show the former. But either way, it may seem surprising that Glanzberg claims that the coordination account is an indirect metasemantics since it claims that semantic values for supplementives in context are always fixed by what appears to be a single factor: recognizable speaker intentions. However, Glanzberg correctly claims that on my view different *kinds* of intentions can fix the semantic value for *pos* in context. For him, this apparently means that my metasemantics is indirect. Though I am indifferent as to whether my metasemantics is direct or indirect in Glanzberg's sense, in discussing the issue he ends up ascribing to me a version of the coordination account I don't hold. This is important since he ends up arguing that his metasemantics is superior to the version of the coordination account I do not hold and his arguments exploit features of that account that differ from the version of the coordination account I do hold. Hence, I need to consider Glanzberg's case that my metasemantics is indirect.

For gradable adjectives like 'tall', 'heavy' and so on it is plausible that at least in many cases speakers intend that degrees on the relevant scale be "cutoff" points for e.g. being tall in a context, since we can identify heights with degrees on the height scale. In such cases, speakers have a height—a number—as the object of their intentions in using the gradable adjective. Hence, the coordination smoothly applies to such gradable adjectives and the degree assigned to *pos* in context will be the height the speaker recognizably intends. But as noted in King (2014a, b), there are gradable adjectives like 'smart' where it is unlikely that in using them a speaker has as the object of her intention a number—a degree on a scale. My thought was that for such gradable adjectives, speakers have as the object of their intention in using them something that *determines* a degree on a scale. In particular, I thought that in the case of 'smart', a speaker might intend a *kind of person as regards intelligence* as the thing that determines the cutoff for being smart in a given context. Presumably, a kind of person as regards intelligence has a degree of intelligence on the scale for intelligence. So in this case, the degree that is assigned to *pos* in context is determined indirectly. The speaker has as the object of her intention a kind of person as regards intelligence, which in turn determines a degree on the relevant scale. This means that there are at least two kinds of intentions invoked by the coordination account in its application to gradable adjectives: intentions whose objects are degrees on the relevant scale; and intentions whose objects are *kinds of things* that in turn determine degrees on the relevant scale. This is enough to make the coordination account an indirect metasemantics according to Glanzberg, though he seems to equivocate on *why* the coordination account is indirect. At one point, he suggests it is because when the object of a speaker's intention is a kind of thing

in using a gradable adjective, the connection between the object of the intention and the degree is indirect with the former determining the latter:

For the standard, on King's theory, the speaker will intend to count certain kind of people as smart. That in turn fixes a standard for smartness. They may understand that this creates a cut-off, but they may well not know there is a parameter whose value needs to be set. King's metasemantics thus marks a clear contrast between simple cases for demonstratives and cases like the standard. *And for the standard, the metasemantic process King proposes already has some indirectness, as it runs through a kind to reach a degree value.*²⁰

But, after arguing that I must invoke other kinds of intentions as well—something I will dispute below—Glanzberg seems to suggest that what makes the coordination account indirect is that it must posit many different kinds of intentions with different kinds of objects:

What this shows, I believe, is that if we wish, we can offer an intention-based indirect metasemantics for parameters like the standard. *We can describe what fixes the standard as a complex set of communicative intentions, making reference to multiple sorts of contextual factors, that must be integrated in a contextually appropriate way to indirectly fix a standard.* We can, with King, require those intentions to be publicly manifest. The contrast with demonstratives still stands. What we have is an indirect metasemantics.²¹

Bracketing the issue of whether the coordination account has to posit more kinds of intentions than I have to this point, resulting in “a complex set of communicative intentions”, which again I return to below, so far I agree with Glanzberg's exposition of the coordination account. However, Glanzberg next appeals to the fact that examples like his 12 discussed above and experimental evidence shows that statistical properties of salient individuals in a comparison class can fix the semantic value of *pos* in context. Of course, I claimed that in example 12 as Glanzberg states it, the conditions of the coordination are not met and so *pos* gets no semantic value in context resulting in ‘He is tall.’ not expressing anything truth evaluable. This explains Glanzberg's judgment about the case that it is hard to hear the sentence as true. And as we saw, when we alter the example so that the coordination account's conditions *are* met, we judge that the sentence *is* true. So I don't think 12 lends any support to the view that statistical properties of salient individuals in a comparison class fixes the semantic value of *pos* in context. As for the experimental data Glanzberg cites, as far as I can see it shows only that conversational participants in using gradable adjectives are sometimes sensitive to statistical properties of salient individuals in a comparison class. This is obviously a far cry from the claim that statistical properties of salient individuals in a comparison class *fix* the semantic value of *pos* in context. My own view is that the reason speakers are sometimes sensitive to statistical properties of salient individuals in a comparison class is that they know such individuals are salient to other conversational participants and they are trying to figure out what they should reasonably intend if they are to be understood; and hearers are aware that such individuals are salient to the speakers

²⁰P. 16, my emphasis.

²¹P. 17, my emphasis.

and are trying to come to know what speakers intend. That, of course, is completely consistent with the coordination account.

Glanzberg goes on to claim that in order to account for cases like 12, in which he claims the semantic value of *pos* is fixed by statistical properties of salient individuals in a comparison class, and his experimental evidence, I am going to have to claim that sometimes the intention that fixes the semantic value of *pos* is an intention whose object is a group of salient individuals in a comparison class:

But the evidence, both from examples like (12) and the experimental evidence I reviewed in section 1, shows that we can also fix standards by reference to statistical properties of salient individuals in a comparison class. We might have intentions corresponding to those salient individuals, which we might express in terms like *the salient things around here*. So we have at least two comparison class-oriented options for how our intentions might go: intention to fix a standard via a kind, or via the statistical regularities across salient individuals.²²

The thought is that such an intention has as its object a group of salient individuals, which in turn determines a comparison class, which with a bit of implicit statistical reasoning determines a degree on a scale. As we have seen, Glanzberg is wrong to claim that the coordination account must posit these sorts of intention to handle 12 and explain his experimental evidence. We have already seen how the coordination account explains these things without positing such intentions.

That said, I don't want to reject intentions of this sort out of hand. There *might* be some cases where a recognizable intention of this sort *does* fix the semantic value of *pos* in context. Here is a plausible case. Suppose Perez is unfamiliar with basketball but has some idea that being tall is an asset in basketball. I extoll the virtues of NBA basketball and get Perez interested in the game. We end up at an NBA game and are allowed to visit the locker room before the game. Perez has never seen an NBA player and he walks into the locker room to find only three players present. We are talking about NBA players, shooting percentages etc. Perez points to the tallest player in the room and says 'Is he tall?' I recognize that Perez wants to know whether the player is a tall basketball player (as would an ideal hearer who knows the common ground of the conversation). Arguably, the object of Perez' intention in using 'tall' is the comparison class comprising basketball players (or the salient players, who in turn determine the comparison class of basketball players), which in turn fixes a degree on the height scale assigned to *pos*. So though I reject Glanzberg's claim that 12 and his experimental data force the coordination account to posit intentions of this sort, we may want to handle cases like the one just described by means of such intentions.

Glanzberg next claims that because of absolute gradable adjectives like 'open' the coordination account is going to have to posit intentions in their use that are intentions to describe things as *completely open*:

We have also seen, both from examples like (10) and experimental evidence, that lexical content can change the situation, leading to much less context dependence, and little reliance

²²P. 16, emphasis in original.

on either of the options we just considered. If we are to take King's line, that would relate to a different sort of intentions [sic], specific to the lexical meanings of absolute adjectives. We might, for instance, have a tacit intention to describe things as *completely open*. So, we have multiple roles for comparison classes and a role for lexical content as part of our metasemantics, even if we view it as intentional. Each comes with different kinds of intentions, and the intentions can be to varying degrees complex and tacit.²³

But Glanzberg is mistaken here in claiming that the coordination account has to posit things like intentions to describe things as completely open to handle absolute gradable adjectives like 'open'. As we saw above, such absolute gradable adjectives are either not contextually sensitive at all or their context invariant meanings alone suffice to secure semantic values in context for them. Either way, they are not supplementives at all and so the coordination account doesn't apply to them. Clearly then, the coordination account doesn't have to posit *any* sorts of intentions to deal with them since it doesn't deal with them at all.

Next, Glanzberg considers how prior overt discourse in using gradable adjectives is viewed on the coordination account (I begin with the end of the last quotation to provide context):

So, we have multiple roles for comparison classes and a role for lexical content as part of our metasemantics, even if we view it as intentional. Each comes with different kinds of intentions, and the intentions can be to varying degrees complex and tacit.

The same can be said for the role of prior discourse. Now, we might see prior discourse as making parts of our complex communicative intention manifest. And so, we might describe the affects [sic] of discourse on the standard in intentional terms. But still, it changes the options for how we are to integrate salient individuals into our process of determining the standard, as we see with example (13). We also see with example (11) and its contrast with (13) that presuppositions can compete with overt discourse for fixing the standard.²⁴

The first two sentences of the quotation here are misleading, suggesting as they do that the coordination account must posit all kinds of new intentions based on the considerations adduced by Glanzberg. Again, I'll return to this below. But Glanzberg is right that the coordination account views overt prior discourse as in his example 13 as one of many ways speakers make their intentions in using supplementives recognizable and hence satisfy condition 2 of the coordination account. As noted earlier, the coordination account views many objective features of discourse—prior discourse, salient objects, questions under discussion—as devices speakers employ to make their intentions recognizable and devices hearers employ to come to know the intentions of speakers. So here Glanzberg gets the coordination account exactly right. However, the very next remarks about changing the options as to how salient individuals are integrated into fixing the standard and how presuppositions compete with overt discourse in fixing the standard on the coordination account miss the mark. In an example like Glanzberg's 13, as discussed above, the overt prior discourse makes completely clear that in this context the speaker is using 5' as

²³Pps. 16–17, emphasis in original.

²⁴P. 17.

threshold for ‘tall’. Hence both conditions of the coordination account are met and 5’ is secured as the semantic value of *pos* in this context. On the coordination account, the salient individuals simply aren’t relevant to assigning a semantic value in context to *pos* in this case. Hence, there simply is not need to “integrate salient individuals into our process of determining the standard.” Though Glanzberg next contrasts 11 and 13, I think he intended to contrast 12 and 13. Glanzberg talks of 12 and 13 showing how presuppositions and overt discourse *compete* on the coordination account. But the coordination account doesn’t view these factors as competing in the way Glanzberg’s view does. In 13, the overt discourse makes clear that the speaker intends 5’ to be the threshold for ‘tall’. Both conditions of the coordination account are met and so 5’ is the threshold for ‘tall’ in context. In 12, as we saw, as formulated by Glanzberg, in the example no threshold meets both conditions of the coordination account and the sentence ‘He is tall’ does not express a truth evaluable content due to *pos* not receiving a semantic value in context. On my amended version of the example, 5’ meets both conditions of the coordination account and so is the semantic value in context of *pos*. But in these cases on the coordination account there is no “competition” between presuppositions and something else, as Glanzberg suggests there is. As we’ll see, Glanzberg appears to be viewing the coordination account through the lens of a theory like his on which various factors compete and combine in fixing a semantic value in context for *pos*. But the coordination account is not like this. Either an intention uniquely satisfies each of conditions 1 and 2 of the coordination account, in which case it *alone* fixes the semantic value in context of the relevant supplementive; or no intention uniquely satisfies each condition, in which case the supplementive has no semantic value in context.²⁵

Finally, Glanzberg summarizes what the coordination account will have to look like, given the considerations he has adduced and the different intentions he claims the coordination account must posit:

If these are all intentions, or the making of intentions manifest, then we have multiple kinds of intentions or ways to make them manifest. In different contexts, these contribute differently to fixing the standard, and how they contribute can be a matter of context. This is a re-casting in intentional terms of the full case I made above for an indirect metasemantics.

What this shows, I believe, is that if we wish, we can offer an *intention-based indirect metasemantics* for parameters like the standard. We can describe what fixes the standard as a complex set of communicative intentions, making reference to multiple sorts of contextual factors, that must be integrated in a contextually appropriate way to indirectly fix a standard. We can, with King, require those intentions to be publicly manifest. . . . What we have is an indirect metasemantics. In light of the arguments I gave in section 1 and here, I believe that the best way to understand King’s proposal is precisely as an intention-based indirect metasemantics, when it comes to parameters like the standard.²⁶

²⁵There are also cases in which felicitously used supplementives are not assigned a unique semantic value in context but rather a *range of candidate semantic values in context*. See King (2018) for discussion.

²⁶P. 17; italics in original text.

This description of how to understand the coordination account is at best *very* misleading. In the first sentence of the quotation, Glanzberg is alluding to all the new kinds of intentions he claims the coordination account must posit, given the cases he raised. But as we saw, none of the considerations Glanzberg raised forced the coordination account to posit *any* new sorts of intentions. In the case of gradable adjectives in the positive form, all of Glanzberg's cases were explained on the basis of the two kinds of intentions the coordination account posited all along. Either the intention will be that a degree on the relevant scale be the semantic value of *pos* in context; or the intention will have as its object a kind of thing as regards the dimension in question (e.g. intelligence), which in turn determines a degree on the relevant scale. I have accounted for all of Glanzberg's cases and data using only these two intentions, which, again, the coordination account posited all along. I did at one point claim that there *may* be cases for which we would want to add a third kind of intention: an intention that had as its object either salient individuals who are members of a comparison class or the comparison class itself, with the comparison class in turn determining a degree on the relevant scale. But the positing of this sort of intention was not required to handle any of Glanzberg's examples or data. Once he had brought up such an intention, I thought there might be cases that were best handled by them and I gave an example of one. As to whether this version of the coordination account is indirect in Glanzberg's sense, I am not sure. But as I indicated, I have no real interest in this question. One point that is important for Glanzberg is that the intentions involved in the use of gradable adjectives will be different from the intentions involved in the use of demonstratives on the coordination account. I think he is right to say they will be, since intentions regarding comparison classes and kinds of things as regards a dimension will not be in play with demonstratives but are with gradable adjectives.

But the other things Glanzberg says about how best to understand the coordination account in the above passage are incorrect. He has the picture of a bunch of different kinds of intentions combining and contributing to fixing a semantic value for *pos* in context and doing so differently in different contexts. But this is completely foreign to the coordination account. To repeat what was said above, in each case where a unique semantic value in context is secured for *pos*, a *single* intention uniquely satisfies each of conditions 1 and 2 in the coordination account. That intention is either an intention with a degree on a scale as its object, or a kind of thing as regards a dimension or a comparison class. There is no "complex set of communicative intentions, making reference to multiple sorts of contextual factors, that must be integrated in a contextually appropriate way to indirectly fix the standard." There is a *single* intention of one of the three kinds mentioned above that uniquely satisfies each of conditions 1 and 2 in the coordination account. This is important, because as we'll see the arguments Glanzberg gives against the coordination account depend on adopting the version reflected in the above passage from Glanzberg. Clearly, this is not the version of the coordination account I am defending.

The primary argument Glanzberg gives for his objective metasemantics and against the coordination account comes in section 3 of his paper.²⁷ As we'll see, Glanzberg really only argues that if one broad view of communication is correct, we should adopt his metasemantics, whereas if another broad view is correct, we should adopt the coordination account. However, I think there is good reason to think that *neither* of the broad views of communication is correct and so as we'll see I think the dialectic is more complicated than Glanzberg makes it out to be. As I indicated above, I also think Glanzberg gets the coordination account wrong here with the result that he would probably be surprised at how much I agree with the broad view of communication he favors. Of course I don't agree that this broad view favors his metasemantics.

Glanzberg makes a number of preliminary points about tense and gradable adjectives and more generally about lexical and functional categories, most of which I agree with. But the rubber meets the road when Glanzberg identifies some of the factors that he thinks "go into our fixing a standard" in context.²⁸ They are:

1. grammatical competence, which delivers a tacit grasp of the structure of the Degree Phrase of which the gradable adjective is a constituent²⁹
2. our ability to approximately represent magnitudes like *fast*, *tall*, *bright* and so on
3. our ability to tacitly do statistics on comparison classes
4. overt prior discourse
5. lexical meaning in the case of absolute adjectives
6. shared presuppositions

I actually agree that, in a sense, all the factors on this list can "go into fixing a standard" in context, with the exception of 5, since as I have said either absolute adjectives are not contextually sensitive or they are like pure indexicals whose context invariant meanings alone serve to secure semantic values for them in context. What is at issue between me and Glanzberg is *how* these factors "go into fixing a standard" in context. To put my cards on the table right off, I think 1 and 2 are background conditions a speaker must satisfy to form the sorts of intentions that the coordination account claims secure semantic values for *pos* in contexts.³⁰ As for 3, I take this ability to play two roles. First, as we have seen I *may* want

²⁷The section is entitled 'Cognition and Semantics.'

²⁸p. 22.

²⁹The grasp is tacit because ordinary speakers are not explicitly aware that its structure is [_{DegP} [_{Deg} *pos*[_A *Adj*]]].

³⁰Note that I don't mention hearers here because so long as the speaker satisfies these condition, she will be able to fix a semantic value in context regardless of whether her *actual* hearers satisfy these conditions. For an idealized hearer who knows the common ground of the conversation is defined to be competent and to have the properties the common ground attributes to the audience, which in usual cases will include having normal cognitive abilities including sensitivity to statistics of comparison classes. When the common ground does *not* attribute this latter property to the audience, the speaker may have to do more than usual to secure a semantic value for *pos* in context according to the coordination account. See King (2014b) for discussion.

to say that we exploit this ability in fixing the standard by having an intention in using a gradable adjective that has either salient individuals in a comparison class, or the comparison class itself, as its object. On the other hand, this ability may constrain what a reasonable speaker who expects to be understood can intend in using a gradable adjective and determine what an idealized hearer who knows the common ground of the conversation knows about the speaker's intention. Similarly, 4 and 6 may constrain what a reasonable speaker who expects to be understood can intend in using a gradable adjective and determine what an idealized hearer who knows the common ground of the conversation knows about the speaker's intention.

Obviously, then, I do not view all of 1–6 as intentions (how could grammatical competence be an intention?!).³¹ Indeed, I don't view *any* of 1–6 as intentions! But Glanzberg seems to suggest that I do, when in discussing 1–6 above he says:

Now, if we stretch the notion of intention enough, I suppose that is alright to call any of these intentions . . . But always putting things in terms of a stretched notion of intention elides important distinctions. Some aspects of the metasemantics of the standard are familiar ordinary communicative intentions. But some are highly tacit abilities. Some are special purpose, while some are general-purpose. These can serve our communicative intentions without being intentions. That is clear from tacit knowledge of syntax, but it holds for tacit abilities to do statistics too. Fixing a standard is a process that draws on many of these distinct kinds of abilities. The version of the indirect metasemantics I prefer keeps track of these distinctions. It can appeal to intentions specifically, in cases like previous discourse, while appealing just as specifically to our abilities to do statistics or our grammatical competence. I think this is the most revealing way to capture the metasemantics of the standard. Hence, I do not opt for a thoroughly intention-based metasemantics. I thus conclude that the metasemantics of parameters is indirect, and also, that the indirectness is not best described in uniformly intentional terms. Intentions are crucial, but not the only factor.³²

Again, I am not claiming any of 1–6 are intentions and so am not eliding distinctions between intentions, special purpose tacit abilities like 1 and general purpose tacit abilities like 3 as Glanzberg claims in this passage. So the case Glanzberg makes here for his objective metasemantics over the coordination account misses the mark.

Glanzberg goes on to provide a reply to the argument I just responded to on behalf of advocates of the coordination account:

There is a reply from defenders of intention-based metasemantics that we should pause to consider. I have highlighted a variety of cognitive resources we draw on to think about properties expressed by gradable predicates, to structure those thoughts in language, and to fix standards in context [1–6 above]. Not all of these are intentions, and not all of them are fundamentally about communication. But still, we can employ basic, often tacit, cognitive resources to serve a variety of functions. Our perceptual abilities serve our motor abilities when we run to catch a ball, for instance. *Is what we see with the standard genuinely a mix of intentions and non-intentional cognitive factors, or rather a mix of intentions, some of which are based on underlying tacit non-intentional factors?* I have suggested the former

³¹Henceforth, when I talk about 1–6 above, I mean to exclude 5 as irrelevant.

³²p. 23.

gives us a more accurate picture of what happens when we fix the standard, and so is a better metasemantics. But it is open to intentionalists to pursue the latter.³³

Note that this isn't how I responded to Glanzberg's argument above. First, to repeat, in any given case where *pos* is assigned a semantic value in context, I don't claim, as Glanzberg seems to suggest here, that there is a *mix* of intentions. I claim that there is a single, recognizable intention. Second, it is misleading to say that I claim that such an intention is *based* on non-intentional things like 1–6. To repeat what I said above about the roles 1–6 play according to the coordination account in fixing semantic values for *pos* in context, 1 and 2 are background conditions a speaker must satisfy to form the sorts of intentions that the coordination account claims secure semantic values for *pos* in contexts. 3 plays two roles. First, I may exploit this ability in fixing the standard by having an intention in using a gradable adjective that has either salient individuals in a comparison class, or the comparison class itself, as its object. On the other hand, this ability may constrain what a reasonable speaker who expects to be understood can intend in using a gradable adjective and determine what an idealized hearer who knows the common ground of the conversation knows about the speaker's intention. Similarly, 4 and 6 may constrain what a reasonable speaker who expects to be understood can intend in using a gradable adjective and determine what an idealized hearer who knows the common ground of the conversation knows about the speaker's intention. So the response Glanzberg makes to his argument on behalf of the coordination account is not my intended response.

Glanzberg ends his paper by outlining two competing views of the role 1–6 play in communication and argues that if one accepts the first, one should adopt the coordination account and if one accepts the second, one should adopt his metasemantics. Here is the first view:

I have already given my reasons why I prefer the less intentional version [Glanzberg's account]. It offers a more refined account of the facts, that keeps track of more information about how the standard is set. But there is a reason that intentionalists can offer for preferring their option. *One might hold that if any of the non-intentional resources I have mentioned—tacit statistics, magnitude representation systems, grammatical competence, etc. [1–6 above]—are to serve communication, they must do so via communicative intentions. If one holds this view, then the only role for these kinds of resources is to produce complex communicative intentions.* If so, then the intentionalist response [sic] we just considered is obligatory, and nearly a tautology.³⁴

Above, I argued that Glanzberg *hasn't* given reasons for favoring his account.³⁵ Also, I clearly do not hold the view that the only role for 1–6 in communication is to produce complex communicative intentions. I have several times now stated

³³Pps. 23–24, my emphasis. Glanzberg footnotes this passage saying that in a personal communication I said this line of thinking was important. In the personal communication in question, I was actually trying to articulate the line of thinking I articulated above in responding to Glanzberg. I fault myself for not being more clear in the personal communication.

³⁴P. 24, my emphasis.

³⁵See the remarks following the material quoted two quotations back about the coordination account “eliding important distinctions.”

the various roles played by 1–6 in the coordination account and these roles that are not a matter of the production of intentions. Indeed, only 3 is implicated at all in the production of intentions³⁶: I tentatively suggested that sometimes the intention that is relevant to fixing the semantic value of *pos* in context is an intention with either salient members of a comparison class or the comparison class itself as its object. But note that even here, the ability to do statistics on salient individuals in comparison class doesn't *produce* an intention. It uses a comparison class that is the object of an intention to yield a degree on a scale. So really *none* of 1–6 is involved in the production of the intentions the coordination account appeals to. Hence to a large extent I actually *reject* the broad view of communication Glanzberg seems to think I hold and that he thinks favors the coordination account.

Glanzberg contrasts the view of the role of 1–6 in communication just discussed that he wrongly attributes to me with the view he holds:

But I also think that the scaffolding upon which we build our overt thoughts and communicative intentions need not be intentional at all, nor need the cognitive and linguistic abilities that provide the scaffolding underlie any specific communicative intentions. To put it simply, when it comes to lower-level cognitive and linguistic abilities, overlap of ability is enough, and we need not have intentions, shared intentions, or any other Gricean-inspired apparatus of communicative intentions corresponding to these. Such overlap of lower-level abilities is enough to serve communication and communicative intentions, even if it is not itself intentional.³⁷

He illustrates this idea with two cases: processing of speech sounds and syntax:

If each of us, as a speaker, processes speech sounds roughly the same way, communication succeeds. It is not necessary to have any communicative intentions surrounding that ability, beyond the general intention to speak and be understood. Our more general intentions to speak and be understood are fully served by having phonetic and phonological abilities, and others having overlapping ones. If they do, we succeed, if they do not, we may fail. I do not see a role for intentions here . . . As with sounds, overlap seems to me to suffice. If my hearer and I share enough grammar, communication succeeds and our communicative intentions are satisfied; if not, it fails. But no communicative intentions need to attach to functional aspects of grammar. Like the processing speech sounds, our understanding of grammar happens automatically, and either it overlaps across speakers or it does not.³⁸

So Glanzberg claims that in the cases of our ability to process speech sounds and our tacit grasp of grammar, we need have no intentions directed towards these abilities for communication to occur and we need only have these abilities sufficiently overlap for communication to occur. I completely agree with the former claim. In general when normal people are communicating, they will have no intentions regarding their ability to process speech sounds or their tacit grasp

³⁶1 and 2 are background conditions a speaker must satisfy to even form the intentions the coordination account claims secure semantic values for *pos* in contexts. In that sense, they are involved in the production of these intentions. But that is hardly a matter of these abilities playing the role of *producing* these intentions.

³⁷P. 24.

³⁸P. 25.

of grammar. Again, that Glanzberg thought I would disagree with this claim suggests that he really did misunderstand my view. However, the latter claim—that all communication requires so far as knowledge of grammar and speech sound processing is sufficient overlap in these abilities—seems to me false at least in the case of knowledge of grammar. Suppose you and I *do* sufficiently overlap in our ability to process speech sounds and our tacit grasp of grammar to communicate. Suppose, however, that I perversely don't believe that your grammar and mine overlap at all. I hear what seem to be English sentences coming from you, but believe that you are speaking a completely different language whose sentences sound like English sentences. Then clearly communication will not succeed, as I will think I have no idea what you are saying. But once we add that what is required for communication as far as this knowledge of grammar goes is sufficient overlap and (tacit) belief that there is sufficient overlap, I am onboard with this claim too.

Finally, Glanzberg returns to the roles of 1–6 above:

The case I made above is that setting the standard relies on these kinds of abilities, and also other abilities like our ability to do tacit statistics, which are not fundamentally communicative at all. That is what I argued for at length. I now add the claim that these can affect communication by merely overlapping across speakers, and we need not, and should not, look for specific communicative intentions to which they are linked. I have not fully supported this claim, by any means. Even so, I hope to have made clear what the two general views of how lower-level cognitive abilities support communication are, and I have offered a few illustrative considerations about why one might prefer my non-intentional option.³⁹

As we saw above, Glanzberg is wrong to think that I claim that 1–6 are “linked” to communicative intentions in any significant way. I don't hold that 1–6 play a role in producing complex communicative intentions as he claims. 1 and 2 are conditions a speaker must satisfy to have communicative intentions at all; and 3–6 can play a role in constraining the kinds of intentions a speaker interested in being understood can have and in determining what an idealized hearer who knows the common ground of the conversation knows about the speaker's intentions. Further, I agree with Glanzberg that with respect to 1–6, communication requires that we overlap sufficiently with respect to these abilities and at least tacitly believe that we do so.

So in the end, the alleged differences between Glanzberg and me in the roles 1–6 play in communication that he has cited so far aren't real differences at all. We pretty much agree on all these points. So far as I can see, our disagreement isn't on some big picture issues about communication as Glanzberg suggests. It is much more local. I claim 1–6 play a role in constraining reasonable speakers' intentions which then determine a semantic value in context for *pos* in favorable cases and determining what an idealized hearer who knows the common ground knows about the speaker's intentions. Glanzberg's claim that 1–6 somehow together and in different ways in different cases *themselves* determine the semantic value for *pos* context. I think the only way to adjudicate *this* disagreement is the old-fashioned way: a case by case consideration of predictions both accounts make in

³⁹Pps. 25–26.

particular examples and a determination of which theory makes the best predictions. Unfortunately, this way of proceeding is currently not an option because, as Glanzberg himself notes at the end of his paper, his metasemantics is not sufficiently fleshed out to make any real predictions. In particular, the exact factors that together fix the semantic value of *pos* in context have not yet been fully identified. Perhaps more importantly, no theory has been offered as to how those factors are weighed and combined in specific cases to fix a semantic value in context for *pos*. Further, Glanzberg claims that the relevant factors are weighed and combined differently in different cases, and no theory has been offered as to what determines how they are weighed and combined in a given case. Until these gaps in the theory are filled in, Glanzberg's objective metasemantics remains more promise than theory.

Let me close with a methodological point that about the dialectic here. Glanzberg and I agree that the coordination account is needed. As I mentioned at the outset, Glanzberg allows that it is the correct metasemantics for what he calls *thematic parameters*. But since we need it anyway, good methodology requires applying it as widely as possible. In other words, we shouldn't posit a second metasemantics until we see that it is absolutely needed in addition to the coordination account. Thus, the argumentative burden is clearly on Glanzberg to demonstrate that his account is needed over and above the coordination account. This is important, because in some of the cases we considered, we seemed to reach stalemates between the competing explanations of the coordination account and Glanzberg's objective metasemantics. But given the dialectic just described, these are wins for the coordination account since a stalemate between the coordination account and Glanzberg's objective metasemantics cannot show the need for his account in addition to the coordination account. But this is precisely what Glanzberg needs to show.

References

- Glanzberg, M. (2007). Context, content, and relativism. *Philosophical Studies*, 136, 1–29.
- Glanzberg, M. (2016). *Not all contextual parameters are alike* (unpublished manuscript).
- Glanzberg, M. (2019). *Indirectness and intentions in metasemantics* (in the present volume).
- Kennedy, C. (2007). Vagueness and grammar: The semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy*, 30, 1–45.
- Kennedy, C., & McNally, L. (2005). Scale structure, degree modification, and the semantics of gradable predicates. *Language*, 81, 345–381.
- King, J. C. (2013). Supplementives, the coordination account and conflicting intentions. *Philosophical Perspectives*, 27(1), 288–311.
- King, J. C. (2014a). The metasemantics of contextual sensitivity. In A. Burgess & B. Sherman (Eds.), *Metasemantics: New essays on the foundations of meaning* (pp. 97–118). New York: Oxford University Press.
- King, J. C. (2014b). Speaker intentions in context. *Noûs*, 48(2), 219–237.
- King, J. C. (2018). Strong contextual felicity and felicitous underspecification. *Philosophy and Phenomenological Research*, 97(3), 631–657.
- Stalnaker, R. C. (1978). *Assertion, syntax and semantics* 9. New York: Academic.

Speakers, Hearers and Demonstrative Reference



Palle Leth

Abstract In the debate on demonstrative reference it is taken for granted that there is such a thing as the semantic instance of determinacy for demonstratives. I argue that the interpretive interaction between the speaker and the hearer suggests that the notion of objective semantic reference in the case of demonstratives is dispensable. Either the speaker and the hearer do not have recourse to any such notion or, at least, there are no reasons for them to have recourse to any such notion. Looking at reactions and interactions does not of course settle the issue whether there is objective demonstrative reference. But it strongly suggests that the issue rather be dismissed.

Keywords Demonstratives · Reference · Intentionalism · Anti-intentionalism · Semantics · Pragmatics · Interpretive interaction

1 Introduction

A demonstrative like *this* or *that* is a linguistic device to refer to a particular object, but the demonstrative itself does not tell us which object it stands for. Augustine says of pronouns among which he includes demonstratives:

[The pronoun] marks out a thing with a less complete signification than a name does. (Augustine 1995, p. 111; *DM* v.13.)

The linguistic meaning of demonstratives need not, however, be conceived of as meagre only. Kaplan thinks of their meaning, or what he calls character, as instructive, in so far as it

provides a rule which determines the referent in terms of certain aspects of the context. (Kaplan 1977, p. 490.)

P. Leth (✉)

Department of Philosophy, Stockholm University, Stockholm, Sweden
e-mail: palle. leth@philosophy.su.se

The question as to which these contextual aspects are has given rise to an intense debate within philosophy of language and theoretical linguistics. There is a plethora of positions and accounts, but at present no consensus.

While it certainly is an important task to discern which factors are relevant to demonstrative reference, the presupposition underlying the debate on demonstrative reference according to which there is such a thing as the semantic instance of determinacy for demonstratives might well be questioned. I will argue that the interpretive interaction between the speaker and the hearer suggests that the notion of objective semantic reference in the case of demonstratives is dispensable.¹ Before I set out to outline the options at stake in our practice with demonstratives, I will briefly consider the recent discussion on demonstrative reference.

2 Intentionalism vs. Anti-Intentionalism

The diverse positions with respect to which factors go into determining the referents of demonstratives can be divided into two basic varieties, intentionalist or anti-intentionalist. Intentionalist accounts take the speaker's perspective in giving precedence to the speaker's referential intention; anti-intentionalist accounts take the hearer's perspective in giving precedence to the cues which are available to the hearer. I will now succinctly describe these two basic approaches and then have a closer look at the typical arguments used in favour of the positions.

2.1 Intentionalist Theories

A starting point for speaker oriented theorists is the idea that the speaker's use of a demonstrative is motivated by her communicative intention to say something about a particular object. Hence, the speaker's referential intention ought to be determinative. If the speaker forms the intention to refer to a particular object in her vicinity, utters *this* and accompanies this utterance by a demonstration, for example a pointing gesture towards the object in question, it is her intention which determines the referent of *this*. Kaplan (in his most recent account) clearly stresses the merely instrumental role of the demonstration:

[T]he demonstration [...] is there only to help convey an intention and plays no *semantical* role at all. (Kaplan 1989, pp. 583–4.)

¹Views similar to those here presented on demonstrative reference have been put forward in Smit 2012, Heck 2014, Gauker 2015, Åkerman 2015 and Bach 2017. However, I hope to present a fuller picture in dispensing with objective demonstrative reference across the board, in taking into consideration the complete variety of interpretive options and in grounding the account in the conversational interaction between the speaker and the hearer.

A fairly extreme version of such an intentionalist approach is the following:

[I]ndexical expressions must be evaluated with respect to what I have called the intended context of interpretation, i.e., with respect to a context providing the co-ordinates intended by the speaker as semantically relevant. (Predelli 1998, p. 113.)

Can the referent of a demonstrative really be just anything the speaker happens to intend to refer to? If actual demonstrations play no role at all, many think that the very possibility of communication is threatened. The extreme intentionalist could argue, though, that the issue of what determines demonstrative reference should be separated from the issue of the possibility of communicative success.

However, the majority of theorists think that the speaker's intention is subject to some form of constraints. Many let constraints operate on the very formation of intentions. Even though the referent of a demonstrative is settled by the speaker's intention, the formation of her intentions is constrained by the general metaphysical principle that one cannot intend what is impossible (cf. e.g. Donnellan 1968 and Davidson 1986). The speaker's use of a demonstrative will refer to the object she has in mind only if the speaker has good reasons to believe that the hearer is in a position to identify the object in question as the referent of the demonstrative. A demonstration is often effective for satisfying this condition. Here is a general statement of this view:

What *U* [utterer] meant by uttering *X* [utterance] is determined solely by *U*'s communicative intentions; but of course the *formation* of genuine communicative intentions by *U* is constrained by *U*'s expectations: *U* cannot be said to utter *X* *M*-intending [*'M*-intention being (roughly) an intention to get an audience to believe (or do) something on the basis of recognising that he or she is so-intended'] *A* [audience] to ϕ [e.g. *M*-intending *A* to believe that *p*] if *U* thinks that there is very little or no hope that *U*'s production of *X* will result in *A* ϕ -ing. (Neale 1992, pp. 552; cf. Neale 2016.)

By this constraint on intention formation, intentionalism may account for the possibility of communication. It should be noted, however, that reference is not affected by the hearer's actual uptake. The hearer may fail to grasp the speaker's intended referent, either because of her incompetence or inattentiveness or for less blameworthy reasons, but as long as the speaker has a reasonable expectation of getting understood, the intended referent counts as the referent. In this context, speaker oriented theorists invoke a distinction between the metaphysical level of reference and the epistemic level of reference. Demonstrative reference is determined at the metaphysical level by the speaker's reasonably formed intention. The epistemic level concerns the hearer's evidential relation to the referent, which is irrelevant for reference determination, just as the way things are taken to be is irrelevant for the way things are (cf. Neale 2016 and Speaks 2017).

Sometimes more importance is given to the hearer's perspective within an intentionalist approach. It may be stipulated that the referent of a demonstrative is the object the speaker intends to refer to, provided that this object also is the object which is most reasonably taken to be the referent by the hearer. King labels this view the coordination account:

[T]he value of a use of a demonstrative in a context is that object *o* that meets the following two conditions: (1) the speaker intends *o* to be the value; and (2) a competent, attentive, reasonable hearer would take *o* to be the object that the speaker intends to be the value. (King 2014, p. 225.)

The distinctiveness of this view appears if we consider that the second condition has to do with what an ideal hearer actually would take the referent to be and also that this view entails that in case the object most reasonably taken to be the referent does not coincide with the speaker's intended object, the demonstrative will lack a referent.

Let us finally note that many intentionalist theorists insist on the plurality of intentions associated with the use of demonstratives. The impetus comes from cases where the speaker intends to refer to a particular object but where the demonstration accompanying the utterance of the demonstrative picks out a different object. Cases like this seem to speak against intentionalist theories, for intuitively the referent is not determined by the speaker's intention in such a case. The conflict, however, is only apparently between the speaker's intention and her demonstration, according to intentionalists. They hold that even though the speaker intends to refer to the object she has in mind, she also intends to refer to the object she is pointing at (Bach 1987, Perry 2009 and King 2013). This latter intention is labelled proximal or directing:

Directing intentions are determinative. That is, the speaker has referred to whoever or whatever plays the role involved in his directing intention. The further intentions are not determinative. Suppose the speaker thinks he is talking to Elwood Fritchey, but is instead talking to Elwood's twin Christopher. When he says 'you' he has referred to Christopher, not Elwood. (Korta and Perry 2010, p. 179.)

The question remains though why it is precisely the directing intention which is determinative.

2.2 *Anti-Intentionalist Theories*

For anti-intentionalists the basic starting point is that the referent of a demonstrative is determined by factors which are available to the hearer, i.e. objective and public features of the context.

What the speaker says, strictly speaking, and thus his reference, is determined by the cues available to the addressee. (Wettstein 1984, p. 74.)

The referent of a demonstrative is the object which is most reasonably taken by the hearer to be the referent. Hence a demonstrative is not invested with a referent by the speaker, but acquires a referent as the result of its insertion into a particular context. On this view, the speaker may end up having demonstratively referred to an object which she did not intend to refer to or which she was not even aware of.

What are these publicly available cues? Some theorists believe that the referent is settled by the demonstration accompanying the demonstrative (Kaplan 1977). The

demonstration may be a pointing gesture and the referent will depend on its spatio-temporal setting:

[T]he referent of a token of ‘that *F*’ is to be the first *F* to intersect the line projected from the pointing finger, i.e. the *F* at the place indicated – one might almost say geometrically – by the accompanying gesture. (McGinn 1981, p. 163.)

But reference may be settled in less precise fashions, e.g. by salience:

For various reasons, an object other than the demonstratum can be more salient. When it is, the salient object rather than the demonstratum is the referent. Neither demonstrations nor speaker intentions are significant in and of themselves; they only affect the reference of demonstratives to the extent that they influence which object is salient. (Mount 2008, pp. 152.)

Other theorists go further and claim that demonstrative reference is settled by all-things-considered judgments:

We can say that the referent of a demonstrative is the object of an all-things-considered judgment, in which the things to be considered are the accessibility criteria, and what has to be determined is which object, if any, adequately and best satisfies those criteria. (Gauker 2008, p. 366; see also Wettstein 1984.)

Some theories of demonstrative reference combine both intentionalist and anti-intentionalist features. Many pointings are not very precise and for some anti-intentionalists it is natural to let the speaker’s intention determine the actual referent within the general direction of her pointing. Even though there are many positions and problems which I have not considered in this brief overview, I think this sketch of the outlines of the debate is sufficient for our present purposes.

2.3 *Method and Argument*

On what grounds are we expected to take sides in this debate? Theorists concoct scenarios where the speaker’s intended referent and the hearer’s assigned referent diverge and try to elicit our intuitions concerning the referent. Here is a case invented by Gauker:

Suppose that Harry and Sally are at a department store and Harry is trying on ties. Harry has wrapped a garish pink-and-green tie around his neck and is looking at himself in a mirror. Sally is standing next to the mirror gazing toward the tie around Harry’s neck and says, ‘That matches your new jacket.’ As a matter of fact, Sally has been contemplating in thought the tie that Harry tried on two ties back. At first she thought she did not like it, but then it occurred to her that it would look good with Harry’s new jacket. We can even suppose that in saying ‘that’ what she intended to refer to was the tie two ties back. But under the circumstances, Harry is in no position to realize that the tie she intended to refer to was the tie two ties back and therefore is in no position to take Sally’s intention into account in identifying the reference of her demonstrative ‘that’. The only thing one could reasonably expect Sally’s demonstrative ‘that’ to refer to is the pink-and-green tie around Harry’s neck. So if we said that the referent of Sally’s demonstrative was the tie she intended to refer to, we could not maintain that the proposition her utterance expressed was a proposition that Harry could assign to it using a method of interpretation that he could reliably employ on

the basis of features of the situation that he could normally be aware of. Instead, we should say that the reference of her demonstrative was the pink-and-green tie around Harry's neck. (Gauker 2008, p. 363.)

According to Gauker's intuitions the speaker's use of *that* refers to the tie the hearer spontaneously takes as the referent. Montminy's intuitions, however, go in the opposite direction:

The intuition invoked by Gauker thus concerns the interpretation Harry may legitimately take to be the correct one rather than the correct interpretation [...] Suppose that later on in the day, Sally and Harry are having this conversation:

Harry: I decided to buy the pink-and-green tie because you said it matches my new jacket.
Sally: I never said that. I was talking about the yellow tie, which you tried on two ties before the pink-and-green one.

It would be odd for Harry to reply, 'Well, I now understand that you were trying to say that the yellow tie matches my new jacket; but what you actually said was that the pink-and-green tie matches my new jacket.' In other words, it would seem unreasonable for Harry to insist that when Sally said 'That matches your new jacket,' 'that' actually referred to the pink-and-green tie, despite Sally's intention to refer to the yellow one. The right thing for Harry to do is to concede that he misunderstood Sally's assertion, even though, he may add, he was quite justified in believing that her utterance concerned the pink-and-green tie. [...] However, instead of saying that Sally's utterance of 'that' failed to refer to what she intended to refer to, we should say, more plausibly, that Sally failed to make clear what her utterance actually referred to. (Montminy 2010, p. 2912–3.)

For Montminy it would simply be odd to hold that the speaker's use of *that* referred to a tie she did not intend to refer to. The fact that one and the same scenario may give rise to opposing intuitions shows the vulnerability of this methodology.

Intuitions must be supported by quite general assumptions about the nature of communication, language and meaning. Gauker draws on epistemic reasons for taking the hearer's perspective:

My objection to [the intentionalist] theory is that it renders the reference of demonstratives inaccessible to hearers. In order to identify the referent of a demonstrative the hearer will have to figure out what the speaker intended to refer to. But apart from an independent interpretation of the speaker's words, hearers will typically be in no position to do that. (Gauker 2008, p. 363.)

It is far from evident however that an intentionalist theory is committed to mind reading of a problematic kind. The speaker's intention as such may be inaccessible to the hearer, but that does not prevent the speaker from providing the hearer access to it by the contextual means at her disposal. Besides, semantic determinacy is one thing, actual communicative success another.

Anti-intentionalists may also invoke the supposed nature of language:

The (or at least a) primary purpose of natural language is to allow for communication concerning the items speakers have in mind and about which they wish to inform others, ask questions, and so on. Natural languages, however, like other *institutions*, e.g. the law, provide for the fulfilment of the institution's primary functions by means of a complex system of rules and conventions. The institutional rules and conventions, although their *point* is to facilitate communication, attain a life of their own once instituted. (Wettstein 1984, p. 65.)

Language as an institution is subject to the separation between execution and result. The speaker has to abide by the rules of the system in order to promote her intention, but precisely because of these rules effects may be generated which are out of the speaker's control. The fact that the language system may run counter to the speaker's interest may, however, equally be a reason for not locating the instance of semantic determinacy there.

Other anti-intentionalists consider the nature of communication:

The presupposition that the theory of reference must recapitulate that of thought seems to have its source in the idea that the purpose of speech is the expression or externalisation of thought; a demonstrative utterance needs to be backed by a demonstrative thought or else it is null and void. [...] The purpose of speech is communication – to get the hearer to have thoughts. And clearly a demonstrative utterance could serve the end of bringing a demonstrative thought to the mind of the hearer, by directing his attention toward some object, without the speaker having the thought the hearer acquires. (McGinn 1981, p. 167.)

If communication is not a matter of the thoughts the speaker has, but of those the hearer comes to have, then perhaps the speaker's intention is not decisive for reference. But such a conception of communication seems certainly quite one-sided.

Here is what an intentionalist says in favour of taking the speaker's perspective:

And I think the concept of locutionary content, in the case of demonstratives, should be tied to directing intentions even in cases, like our second case, where what is said seems, on balance, to go with demonstrations. We need a concept very close to what is said, but one that seals off the forensic issues. [...] A prima facie plausible principle is that if one is sincere, and semantically competent, and makes no verbal slips, what one says will have the same content as the belief one intends to express by saying it. (Perry 2009, pp. 192–3.)

Perry thinks that our ordinary concept of what is said is permeated by extraneous issues such as epistemic access and accountability, and therefore it had better be replaced by the novel concept of locutionary content. This concept should clearly reflect the structure of the speaker's thought and belief. But this seems to beg the question. Is it not equally plausible that if one is semantically competent, the interpretation one comes up with is what is said by the speaker. On what grounds does Perry take the speaker's faultlessness to be regulative?

Neither intuitions concerning specific cases nor considerations of a more general nature yield any decisive arguments concerning demonstrative reference. It is plausible that a demonstrative refers to what the speaker intended to refer to, but equally plausible that it refers to what the hearer is in a position to take it to refer to. The presupposition common to both intentionalists and anti-intentionalists is that in addition to the speaker's intended referent and the hearer's assigned referent, there is the actual, objectively correct, referent. Perry very explicitly asks:

We are pretty clear about what [the speaker] intended to say, and what [she] was understood as saying. But what did [she] say? (Perry 2009, p. 191.)

The alleged necessity to pass a verdict also shows up in the notion of counting the speaker as having referred:

Intuitively, a speaker succeeds in referring if she has arranged things in such a way that the audience can reasonably be expected to recognize her intention. If for some abnormal reason they do not, the speaker still counts as having referred. (Stokke 2010, p. 390.)

Intentionalists are concerned with the conditions for the speaker's counting as having referred. It is one thing, however, to state what the speaker reasonably must do in order to realize her intention to refer. It is another thing to stipulate that the speaker's doing what she is reasonably expected to do amounts to the accomplishment of reference. If the speaker's reference is not heard and understood, what is the use of counting her as having referred? (What does, by the way, counting her as having referred amount to, if not her being heard and understood by the theorist?)² If the question of what the actual referent is is not only hard to settle, but also does not serve any particular purpose, we had perhaps better not ask it. This is in any case what an inquiry into interpretive interaction will suggest.

3 The Evidence from Interpretive Interaction

Theorists interrupt their scenarios at the juncture of non-convergence between the speaker's intention and the hearer's understanding and solicit us to pass a verdict. However, if we look at the interpretive interaction which follows upon the problematic utterance, the conviction that a verdict should be passed will certainly appear less motivated.

Let us consider Kaplan's famous case involving pointing and a replacement of pictures:

Suppose that without turning and looking I point to the place on my wall which has long been occupied by a picture of Rudolf Carnap and I say:

(1) [That] is a picture of one of the greatest philosophers of the twentieth century.

But unbeknownst to me, someone has replaced my picture of Carnap with one of Spiro Agnew. (Kaplan 1978, p. 30; adapted.)

I propose to consider the likely directions of interpretive interaction between the speaker and the hearer once (1) has been made.

²Cf. Reinach 1913, p. 19 [ch. I, § 3]: 'The command is according to its essence *in need of being heard* [*vernehmungsbedürftig*, i.e. in need of being apprehended, recognized or received – taken up]. It can of course happen that commands are given without being heard. Then they fail to fulfil their purpose. They are like thrown spears which fall to the ground without hitting their target.' It seems to me that many theorists who insist that reference occurs provided that certain conditions on the speaker's part are fulfilled unnecessarily suggest that the speaker's attempt at referring somehow hits the target without being taken up.

3.1 *Conflicting Evidence*

Let us first suppose that the hearer is aware that the picture on the wall depicts Spiro Agnew and also knows that Agnew was a politician. Then she will certainly be a little puzzled about the utterance. On the one hand, the absurdity of the predication makes it unlikely that the speaker wants to refer to Agnew. On the other hand, the speaker does point at the picture of Agnew. The evidence goes in opposing directions.

It is not plausible that the hearer in this situation would ask herself what the actual referent of *that* in (1) is and attempt to settle this question one way or other. The hearer would rather ask herself whether there has been a mistake or whether the speaker is joking, and then ask the speaker what she intended to refer to. When the speaker becomes aware of the replacement of pictures, she will refer to Carnap in a way which makes evident to the hearer that the predication is concerned with him.³

In the case of this novel utterance, there is communicative success in so far as it is understood by the hearer to the speaker's satisfaction. In order to explain this, it is sufficient to assume that the hearer's assigned referent corresponds to the speaker's intended referent. We do not have to assume that the hearer's assigned referent and the speaker's intended referent correspond to any such thing as the actual referent (cf. Neale 2007).

What are we to say about the demonstrative reference in (1)? It seems adequate to say that (1) is erased and replaced by the speaker's novel utterance. The question as to which the referent of *that* is in (1) is neither settled nor asked.

3.2 *Replacement*

Let us now suppose, as above, that the hearer knows that the picture depicts Spiro Agnew and that she also knows that Agnew was a politician. But unlike the previous case, the hearer takes (1) to deliver new information about Agnew. She exclaims: 'I did not know Spiro Agnew also was a philosopher!', thinking to herself 'This is incredible, but true, since the speaker tells me so' (and not 'This is too incredible to be true, therefore it cannot be what the speaker wanted to say').

The hearer's exclamation will make the speaker aware that something has gone wrong and when she has discovered the replacement of pictures, she will tell the hearer to whom she wanted to refer. Does the fact that the original utterance was initially taken by the hearer to make a certain claim make a difference as to the issue of the reference in (1)?

³In many cases, of course, the hearer will not be able to ask the speaker what she wanted to refer to. Her interpretive attitude remains the same though: she wants to know what the speaker wanted to refer to, not what was actually referred to.

The hearer assigns a referent to the demonstrative in (1) and she takes (1) to deliver a certain piece of information. Upon the speaker's reaction to her interpretation, the hearer considers what she took (1) to say and what the speaker wanted to say and given that she is informed that the speaker did not want to say what she took her to say, she erases (1), which – for some reason or other – only caused confusion and replaces it, as it were, with the novel utterance.

For some theorists it is now natural to invoke the distinction between what the speaker wanted to refer to and what the speaker managed to refer to. The question the hearer sets out to answer though is simply what the speaker wanted to refer to. Nothing in the hearer's procedure corresponds to her being concerned with what the speaker actually referred to. There were certainly reasons for the hearer to take the referent to be the picture of Agnew, but there are no reasons for neither the hearer nor the speaker to establish any such thing as the actual referent of the demonstrative in (1). This is not to say that there could be no such reasons, only that this is not a case where such reasons present themselves.

We might of course imagine that the hearer's attitude amounts to the view that the speaker actually referred to the picture of Agnew, but because the speaker says that she wanted to refer to a picture of Carnap, the hearer disregards the actual reference of the demonstrative. The hearer's taking an interest in the speaker's intended referent certainly does not prevent the hearer from thinking that there is such a thing as the objectively correct referent. But if the hearer is prepared to erase the original demonstrative, objective reference does not become an issue between her and the speaker and any convictions the hearer (or the speaker) has in this regard remain at the level of a merely subjective conviction. The fact that erasure and replacement do not exclude that there be such a thing as the objective referent is without consequence as long as the hearer acts as if there were the speaker's intended referent only.

In the two cases considered so far there is non-convergence between the speaker's intended referent and the hearer's interpretation, but this non-convergence does not give rise to disagreement as to the interpretation of (1), but to the erasure of (1) and the replacement of it by a novel utterance. I propose to describe the hearer's interpretive attitude in cases like these as opting for the speaker's intended meaning. As far as I can see this interpretive attitude has no implications for any such thing as the actual referent.

3.3 *A Case of Consequence*

Let us now consider a slightly modified version of Kaplan's Carnap vs. Agnew case:

Suppose that without turning and looking the speaker points to the place on her wall which has long been occupied by a picture of Rudolf Carnap and says:

(2) That is your birthday present.

But unbeknownst to the speaker, someone has replaced the speaker's picture of Carnap with a valuable painting of Spiro Agnew.

Upon the utterance of (2), we can imagine the following exchange taking place between the speaker and the hearer:

Hearer Oh, thank you very, very much! I am really most grateful.

Speaker [*aside*] How can she get so grateful about that Carnap picture? [*turning around*] Oh, I am sorry. There must have been an unfortunate replacement of pictures. I did not refer to the Agnew painting. I referred to the picture of Carnap which, as you certainly know, has long occupied my wall.

Hearer I understand that you did not want to refer to the Agnew painting, but that is nevertheless what you actually did refer to, that painting is what was referred to by the demonstrative in your utterance together with your pointing gesture, and now that painting is mine.

Speaker Perhaps my use of *that* did not refer to the Carnap picture, as you seem to think, but for all that it did not refer to the Agnew painting, since my use of a demonstrative cannot refer to something which I do not intend to refer to.

Hearer The reference of a demonstrative is not settled by the intention with which it is uttered, but by the public and accessible features of the context in which the demonstrative is uttered: you clearly pointed at the Agnew painting. All things considered, how could I know you did not want to give me the Agnew painting? Besides, since you seem to take your intentions so seriously, did you not have, in addition to the intention to refer to the Carnap picture in your mind, the intention to refer to the picture on the wall behind you? The Agnew painting is mine in virtue of your directing intention, as they call it.

Speaker You know very well I am not in the habit of presenting so valuable gifts to my friends.

The hearer could of course erase the careless utterance (2) and replace it with a novel utterance in line with the speaker's eventually declared intention, as in the cases we considered above. The hearer may also, however, refrain from doing so. Our interest in speaker intentions has its limits and sometimes charity is not a transcendental condition of understanding. The hearer may refuse to erase the original utterance and enforce her own interpretation of it instead. The hearer and the speaker each cling to their own preferred referent and invoke diverging principles of interpretation, so that a serious disagreement as to what is the referent of (2) is brought about. Did the speaker by uttering (2) and by pointing the way she did actually refer to the Agnew painting? A theory of demonstrative reference is perhaps called for in order to adjudicate between the speaker's and the hearer's competing claims. Actual reference seems to become an issue for the speaker and the hearer and possibly also for the theorist.

Even though the hearer seems to be making a claim as to the actual referent of the demonstrative, it seems to me that by invoking objective reference the hearer in our imagined scenario is simply emphasizing the fact that she had good reasons to take the demonstrative as referring to the Agnew painting. The necessity of objective reference is apparent only. Construing her claim as concerned with merely

the reasonableness of her interpretation will actually serve the hearer's purposes better. Whereas any claim as to the objectively correct referent will be ineffective against the speaker in so far as it will easily be disputed, the speaker cannot but concede that the referent of *that* was not unreasonably taken by the hearer to be the Agnew painting. The speaker can hardly deny that she did not succeed in making her intended referent manifest to the hearer. And she cannot invalidate the hearer's interpretation simply because it does not correspond to her intended referent. The issue is what the referent was most reasonably taken to be irrespective of the speaker's actual intention.

This does not imply that whatever the hearer takes as the referent can be held against the speaker. In order for the hearer to be able to enforce her interpretation, it must be reasonable and even be the most reasonable interpretation given the publicly available cues. The speaker will certainly not want to be held responsible for anything less than the most reasonable interpretation of her utterance. Ultimately then the hearer must be able to claim that any hearer in her position would have taken the referent to be what she took it to be. The relativization to the hearer's position is reasonable, I think, for the reason that the speaker addressed herself to the hearer, which implies taking the hearer's prospects of understanding into account.

To a certain extent, the prospect of agreement between the speaker and the hearer is enhanced if they stop arguing about what the objectively correct referent of the demonstrative is and confine themselves to the two separate issues of what the referent was most reasonably taken to be and what it was intended to be. Any conviction as to what constitutes the correct referent, whether it is the intended referent or the most reasonable referent, had better be put on the side. An argument as to the correct referent must make use of special semantical and metaphysical principles the content and application of which are not straightforward even for professionals, whereas an argument as to the most reasonable interpretation only makes use of the kind of reasoning which occurs in ordinary disputes. The prospect of agreement should not be exaggerated though, for what the most reasonable interpretation is may not be easily settled. Abandoning the notion of objective reference may well cast us into the fire.

It would be a mistake however to think that a theory of demonstrative reference could save us from this predicament. For in the cases where the hearer does not opt for the speaker's intended referent, but enforces her own referent, it is precisely, despite what she may sometimes profess, the mere reasonableness of her interpretation which is at stake. The hearer would be unaffected by any theory telling her what the objectively correct referent is, for irrespective of what any such thing as the correct referent happens to be, she had the best reasons to take the referent to be what she took it to be and her claim is concerned with this issue only. Therefore, only arguments as to reasonableness can be relevant. The verdict of a reasonable judge (cf. Travis 1989) is called for, both for establishing what the most reasonable interpretation is and for laying down what the reasonable consequences of the reasonable interpretation should be.

I propose to describe the hearer's interpretive attitude in cases of disagreement as opting for the most reasonable referent. Intentionalists will insist that the most

reasonable interpretation does not represent the real reference. The metaphysical instance of determinacy is the speaker's intention, the most reasonable interpretation is of epistemic concern only and does not amount to the correct interpretation. Anti-intentionalists on the contrary will say that disagreement proves that it is precisely the most reasonable referent which constitutes the objectively correct referent. Disagreement is indeed a candidate case for the possible relevance of actual reference. As I see it, however, the dispute turns out to be a matter of what is the most reasonable reference. Consequently, again, there seems to be no reason to say that anything constitutes the objectively correct referent.

3.4 Merriment

Consider now the original Carnap vs. Agnew case again. In response to (1), the hearer may well burst out laughing:

Spiro Agnew, the greatest philosopher of the 20th century!

When the speaker is made aware of which the picture behind her is, she may well laugh too. And when they both have laughed for a while, the speaker tells the hearer which person she does think is the greatest philosopher of the twentieth century.

Before the original utterance was erased and replaced by a novel utterance which the hearer understood to the speaker's satisfaction, what were the speaker and the hearer laughing at? Was it that the speaker actually had referred to a picture of Agnew and had made the assertion that Agnew was the greatest philosopher of the twentieth century? It seems to me that the referent is rather imagined to be the picture of Agnew. The hearer and the speaker imagine that the speaker had referred to the picture of Agnew or that someone had referred to it or that the hearer had taken the speaker to refer to it or that someone had taken the speaker to refer to it. They need not have recourse to any notion of the objectively correct referent in order to laugh.

Even if the imagined referent is not the most reasonably assigned referent, it is not unreasonable. There must be some more or less clear connection between the situation at hand and the imagined referent. If the hearer's imagined referent is too far fetched, the merriment will not be mutual or at least not wholehearted. This is the occasion to consider a case invented by Reimer and her account of it:

Suppose, for instance, that I suddenly realize that I have left my keys on the desk in my (shared) office. I return to my office, where I find the desk occupied by my officemate. I then spot my keys, sitting there on the desk, alongside my officemate's keys. I then make a grab for my keys, saying *just as I mistakenly grab my officemate's keys*, "These are mine." Now in such a case, Kaplan would presumably want to say that I had a "directing intention" with regard to *my* keys. For it was *my* set of keys that I focused on, and it was *my* set of keys that "directed" my grabbing. I intended to grab *my* keys, not my officemate's. Thus, on Kaplan's view, the demonstratum of "these," as that expression occurred in my utterance, was *my* set of keys. And yet the keys that I've demonstrated by way of grabbing are my officemate's keys. And so surely in such a case my officemate would speak truly were he to

say to me, “No, you’re wrong. *Those* are not *your* keys; they’re mine.” The appropriateness (not to mention truth) of such a reply would suggest that my officemate’s keys—and not my own—were the demonstratum of the demonstrative expression occurring in my utterance. For if *my* keys were that demonstratum, then my officemate’s allegation that what I had uttered was untrue, would have been false—which surely it is not. If Kaplan’s view were correct, then my officemate’s reply would indicate that he simply hadn’t *understood* what I had in fact said (which was actually *true*), and his failure to understand what I had said, would be due, on Kaplan’s view, to the fact that I demonstrated an object which was not the actual (i.e., intended) demonstratum. (Contrast the appropriateness of my officemate’s reply with the inappropriateness of my rejoinder, uttered while handing him back his keys, “Yes, *these* keys are yours, but I never said they were mine.” If Kaplan’s view were correct, my rejoinder ought to be both appropriate and true.) (Reimer 1991, pp. 190–1.)

Reimer assumes that categorical statements as to what was referred to and what was true and what was false are appropriate. But is it really appropriate for the hearer to take a stance on the falsity of the speaker’s utterance of ‘These are mine’, since she does not believe that the speaker wants to say that the hearer’s keys are hers, the speaker’s? There is rather occasion for repair than for truth evaluation. The appropriate thing for the hearer to do in cases of confusion is not to assert that the speaker spoke falsely, but to ask the speaker what she wanted to say. The hearer’s saying ‘No, you’re wrong’ might not be inappropriate however, since that response need not be concerned with the objective reference, but simply with the imagined reference. The hearer imagines that the speaker had asserted that the hearer’s keys were her keys and she gives the speaker her keys with a smile. It would be inappropriate for the speaker to say ‘I never said they were mine’, because the appropriate thing to say is that she never wanted to say that they were hers. Since no theoretical reasons nor any intuitions provide us with any guidance in cases like this any verdict as to what was actually referred to will be arbitrary and is in any case uncalled for.

4 Conclusion

The cases that I have considered are certainly not exhaustive, but they display hopefully a sufficient amount of variety in order to make plausible that neither the convergence nor the divergence of the speaker’s intention and the hearer’s interpretation requires such a thing as the objectively correct semantic referent of a demonstrative in context. Either the speaker and the hearer do not have recourse to any such notion or, at least, there are no reasons for them to have recourse to any such notion. A rational reconstruction of the interpretive interaction between the speaker and the hearer shows that there are basically three questions at issue between them in cases of demonstrative reference:

- (i) What does the speaker want to refer to?
- (ii) What is the speaker most reasonably taken to refer to?
- (iii) What could the speaker be imagined to refer to?

None of the questions (i), (ii) or (iii) amounts to nor has any implications for the question (iv):

(iv) What does the speaker as a matter of fact refer to?

Nevertheless, it is question (iv) which has been taken by theorists of demonstrative reference to require a determinate answer. I hope to have shown that the context in which demonstrative reference occurs is not a context where correctness is at issue. The three options above are equally legitimate and which option is taken depends fundamentally on the hearer's interest, whether she is interested in the speaker's intention or in holding the speaker responsible for her utterance or in playing with her utterance.

Looking at reactions and interactions does not of course settle the issue whether there is objective demonstrative reference. But it strongly suggests that the issue rather be dismissed. To the extent that a theory of demonstrative reference in context seeks to account for our use of demonstratives it seems that the fact that objective semantic reference is dispensable in practice speaks in favour of its being dispensed with also in theory. If so, we would have, by abandoning this notion, reconciled speaker oriented and hearer oriented theories of demonstrative reference.

Acknowledgement The research reported in this paper was supported by the Swedish Research Council under project no. 437-2014-255.

References

- Åkerman, J. (2010). Communication and indexical reference. *Philosophical Studies*, 149(3), 355–366.
- Åkerman, J. (2015). The communication desideratum and theories of indexical reference. *Mind & Language*, 30(4), 474–499.
- Augustine, (1995). The teacher [DM]. In Peter King (Trans.) *Against the academicians and the teacher*. Indianapolis/Cambridge: Hackett (pp. 94–146).
- Bach, K. (1987). *Thought and reference*. Oxford: Clarendon Press.
- Bach, K. (2017). Reference, intention, and context: Do demonstratives really refer? In M. de Ponte & K. Korta (Eds.), *Reference and representation in thought and language* (pp. 57–72). Oxford: Oxford University Press.
- Davidson, D. (1986). A nice derangement of epitaphs. In *Truth, language, and history* (pp. 89–107). Oxford: Oxford University Press, 2005.
- Donnellan, K. S. (1968). Putting Humpty Dumpty together again. *The Philosophical Review*, 77(2), 203–215.
- Gauker, C. (2008). Zero tolerance for pragmatics. *Synthese*, 165(3), 359–371.
- Gauker, C. (2015). The illusion of semantic reference. In A. Bianchi (Ed.), *On reference* (pp. 21–38). Oxford: Oxford University Press.
- Heck, R. G., Jr. (2014). Semantics and context-dependence: Towards a Strawsonian account. In A. Burgess & B. Sherman (Eds.), *Metasemantics: New essays on the foundations of meaning* (pp. 327–364). Oxford: Oxford University Press, 2015.
- Kaplan, D. (1977). Demonstratives: An essay on the semantics, logic, metaphysics, and epistemology of demonstratives and other Indexicals. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan* (pp. 481–563). New York/Oxford: Oxford University Press, 1989.

- Kaplan, D. (1978). Dthat. In P. Yourgrau (Ed.), *Demonstratives* (pp. 11–33). Oxford: Oxford University Press, 1990.
- Kaplan, D. (1989). Afterthoughts. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan* (pp. 565–614). New York/Oxford: Oxford University Press.
- King, J. C. (2013). Supplementives, the coordination account, and conflicting intentions. *Philosophical Perspectives*, 27(1), 288–311.
- King, J. C. (2014). Speaker intentions in context. *Noûs*, 48(2), 219–237.
- Korta, K., & Perry, J. (2010). Intentions to refer. In L. Baptista & E. Rast (Eds.), *Meaning and context* (pp. 161–186). Bern: Peter Lang.
- McGinn, C. (1981). The mechanism of reference. *Synthese*, 49(2), 157–186.
- Montminy, M. (2010). Context and communication: A defense of intentionalism. *Journal of Pragmatics*, 42(11), 2910–2918.
- Mount, A. (2008). Intentions, gestures, and salience in ordinary and deferred demonstrative reference. *Mind & Language*, 23(2), 145–164.
- Neale, S. (1992). Paul Grice and the philosophy of language. *Linguistics and Philosophy*, 15(5), 509–559.
- Neale, S. (2007). On location. In M. O’Rourke & C. Washington (Eds.), *Situating semantics: Essays on the philosophy of John Perry* (pp. 251–393). Cambridge, MA: The MIT Press.
- Neale, S. (2016). Silent reference. In G. Ostertag (Ed.), *Meanings and other things: Themes from the work of Stephen Schiffer* (pp. 229–342). Oxford: Oxford University Press.
- Perry, J. (2009). Directing intentions. In J. Almog & P. Leonardi (Eds.), *The philosophy of David Kaplan* (pp. 187–201). New York: Oxford University Press.
- Predelli, S. (1998). I am not here now. *Analysis*, 58(2), 107–115.
- Reimer, M. (1991). Demonstratives, demonstrations, and demonstrata. *Philosophical Studies*, 63(2), 187–202.
- Reinach, A. (1913). The Apriori Foundations of the Civil Law. *Aletheia* (Vol. III, 1983, pp. xxxii–xxxiii & 1–142, John F. Crosby, Trans.).
- Smit, J. P. (2012). Why bare demonstratives need not semantically refer. *Canadian Journal of Philosophy*, 42(1), 43–66.
- Speaks, J. (2016). The role of speaker and hearer in the character of demonstratives. *Mind*, 125(498), 301–339.
- Speaks, J. (2017). A puzzle about demonstratives and semantic competence. *Philosophical Studies*, 174(3), 709–734.
- Stokke, A. (2010). Intention-sensitive semantics. *Synthese*, 175(3), 383–404.
- Travis, C. (1989). *The uses of sense: Wittgenstein’s philosophy of language*. Oxford: Oxford University Press, 2001.
- Travis, C. (1997). Pragmatics. In *Occasion-sensitivity: Selected essays* (pp. 109–129). Oxford: Oxford University Press, 2008.
- Wettstein, H. K. (1984). How to bridge the gap between meaning and reference. *Synthese*, 58(1), 63–84.

How to Say When: A Reichenbachian Approach to the Answering Machine Puzzle



Agustin Vicente and Dan Zeman

Abstract In this paper, we offer a novel solution to the much discussed “answering machine puzzle” and similarly problematic cases for the Kaplanian view of temporal indexicals. The solution we propose consists in an appeal to a well-established and (for many still) useful framework: Reichenbach’s theory of tense and aspect. Starting from some more recent articulations of the theory in its application to temporal adverbials, we show how it can be applied to ‘now’ so that to provide an easy and intuitive solution to the puzzle. In addition, the proposed solution, while departing in several important respects from the Kaplanian orthodoxy, remains conservative in that it preserves some of Kaplan’s central tenets about ‘now’, thus making it less radical than many of the alternative solutions found in literature.

Keywords Temporal indexicals · Answering machine puzzle · Kaplan · Reichenbach

1 Introduction

In contemporary philosophy of language there’s an interesting debate as to how to give an account of utterances of ‘I am not here now’, which sound contradictory but make perfect sense when played, for example, by an answering machine. In Kaplan’s (1989) framework, *context* (as opposed to circumstance) plays the role of providing values for indexicals and other context-sensitive expressions. It does so via *character*, which specifies rules that help assign values to indexicals in context, such as:

A. Vicente
Ikerbasque Foundation for Science, Bilbao, Spain

Linguistics Department, University of the Basque Country (UPV/EHU), Leioa, Spain
e-mail: agustin.vicente@ehu.eus

D. Zeman (✉)
Department of Analytic Philosophy, Slovak Academy of Sciences, Bratislava, Slovak Republic

“I” picks up the speaker of the context;
 “now” picks up the time of the context;
 “here” picks up the place of the context; etc.

Thus, for Kaplan, sentence

(0) I am here now

is true if the speaker of the context is at the time and place of the context. But since this seems to happen in every context, (0) is a logical truth. If so (assuming a very simple rule for negation), the negation of (0), namely

(1) I am not here now.

should be false in every context – that is, a contradiction. Yet, when played by an answering machine, (1) is true. This particular problem for the Kaplanian framework is what authors like Sidelle (1991) have dubbed “the answering machine paradox” or, less dramatically, “the answering machine puzzle”.¹

The answering machine puzzle makes clear that Kaplan’s theory has to be revised, but it’s much less clear how it ought to be revised. Things are further complicated by the fact that the problem is not limited to answering machines, but more general. The debate about the answering machine case has given rise to multiple other examples where indexicals are used in a non-standard (i.e. non-Kaplanian) way. For instance, imagine that Anne leaves a post-it note on the desk of a colleague that comes at work only on Thursdays at a certain time containing the sentence

(2) There is a meeting now.

Anne leaves the post-it on her colleague’s desk on Monday, as soon as she finds out that the meeting has been scheduled for Thursday at the time the colleague comes in. Arguably, in such a scenario (2) does not mean that there is a meeting on Monday at the time Anne is writing the note, but, if anything, that there is a meeting on Thursday at the time of her colleague’s arrival. Second, the problem is not limited to ‘now’ intuitively referring to times *after* the original utterances have been made: as Predelli (1998) shows, the so-called “historical present” puts pressure on the Kaplanian theory by having ‘now’ intuitively referring to times *before* the relevant utterances take place – when, for example, a historian narrates Napoleon’s deeds in a certain year by uttering

(3) Napoleon, now commander of the French troops in Italy, defeats the Sardinian forces and turns against Austria.

The ‘now’ here intuitively refers to the time when Napoleon was acting, not to the time of utterance. Third and finally, it is not only ‘now’ that creates problems:

¹We rest content here with presenting the answering machine puzzle in a very informal way. For more formal renderings of the problem, see, e.g., Predelli (1998) or, more recently, Briciu (2018).

‘here’ can also be used to refer to locations other than that of the context of utterance – which can be seen both in the answering machine case and in examples of historical past (Predelli 1998).² In all these cases the problem stems from the fact that the relevant expressions refer to contextual parameters (time, location etc.) whose values are not those of the context of utterance.

The strategies used to deal with these cases are diverse. Some accounts (e.g., Predelli 1998) try to accommodate all the examples, while others argue that not all are equally significant for a semantic theory and thus that they can be safely ignored (Cohen 2013; Michaelson 2013). Among the authors that take all the data into consideration, some account for it pragmatically (e.g. Ruffino 2012), others semantically. The semantic accounts themselves are varied: Predelli (1998) and Åkerman (2009) argue for an intention-based account, whereby the reference of an indexical depends on the intentions of the speaker; Corazza, Fish and Gorvett (2002) appeal to conventional rules, such that different conventions fix the relevant agents, places and times; in a similar vein, Gauker (2008) contends that no intentions are needed, but also that no particular conventional rules determine the reference of indexicals; Romdenh-Romluc (2006) proposes an audience-based approach; Mount (2008, 2015) an account based on coordination between speakers and hearers; Cohen (2013) adopts a “tokening” view, so that the content of indexical elements is related the context of their tokening; Michaelson (2013) opts for a “shifty characters” proposal, according to which characters themselves depend on the communication channel, echoing earlier claims to the effect that indexicals like ‘now’ are ambiguous (Smith 1989); Sidelle (1991) and, more recently, O’Madagain (2014) and Briciu (2018) appeal to “deferred utterances” or other mechanisms for making remote utterances to solve the puzzle.

All these proposals have been found to be problematic in one way or another: they either fall prey to counterexamples, have to bite some (unintuitive) bullet, or are found unpalatable on independent grounds (see Cohen and Michaelson 2013 for an overview). In this paper, our aim is not to engage with the vast literature on this specific puzzle, but rather to put forward yet another solution to it. The solution we propose consists in appeal to a well-established and (for many still) useful framework: Reichenbach’s theory of tense. Starting from some more recent articulations of the theory in its application to temporal adverbials, we show how it can be applied to ‘now’ so that to provide an easy and intuitive solution to the answering machine puzzle and the other problematic cases. What’s more, the solution we propose, while departing in several important respects from the Kaplanian orthodoxy, at the same time preserves some of Kaplan’s central tenets about ‘now’, thus making it less radical than many of the alternative solution found in literature.

²We leave what have been dubbed “demonstrative” uses of ‘here’ aside (as when, pointing to a map, one says ‘I’m here’ – see Smith (1989) for discussion). We also avoid discussion of the possibility of ‘I’ referring to other agents than the speaker. Our main focus is on ‘now’.

We structure the reminder of the paper as follows. In the next section, we present the main features of the Reichenbachian framework we adopt. Drawing on more recent work that uses that framework to treat temporal adverbials, in Sect. 3 we show how it can be applied to solve the problematic cases involving ‘now’, including answering machines and the historical past. Section 4 is an interlude meant to strengthen our case by drawing an analogy between time and space. Finally, in the last section we compare our theory both with the Kaplanian framework (showing that, while it departs from it in important ways, it also preserves some of its essential claims) and with several other (families of) solutions found in literature. We end by listing what we take the advantages of our theory to be.

2 The Framework

It is a well-known fact that languages treat time in terms of a continuous dimension that goes from past to future (or from future to past). One prominent theory of tenses relying on the idea of succession of times is Reichenbach’s (1947).³ Reichenbach distinguished three kinds of temporal parameters that are relevant to account for our talk about time (which includes tense and aspect), which are commonly referred as Utterance Time (UT), Event Time (ET) and Reference Time (RT). UT is the time at which the utterance is made and ET denotes the time when the event occurs. RT is the time the speaker talks about, and its need is illustrated by the interpretation of the perfect (helpful examples will be provided shortly). Klein (1995), who calls it “assertion time”, characterizes it as “the time to which the assertion is confined, for which the speaker makes a statement” (187). Klein also holds that tense (past, present, future) relates RT to UT, while aspect (perfective, progressive) relates ET to RT. Ever since Reichenbach’s (1947), Reichenbachians have assumed that UT relates to ET only via RT.

The postulation of three temporal parameters creates a certain degree of flexibility that allows the framework to handle various temporal constructions. That is, the three temporal parameters can stand in all kinds of relations of succession (or coincidence), thus yielding various temporal structures we call “reference frames”. With simple tenses (past, present or future), the relation between ET and RT is one of *central coincidence* (Demirdache and Uribe-Etxebarria 2007), where ET and RT occupy the same points in the reference frame. Consider, for example, a usual utterance of a simple present sentence like

(4) I am thirsty.

We can graphically represent the reference frame corresponding to (4) as follows:

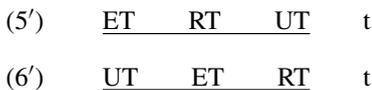
(4') $\underline{UT = RT = ET} \quad t$

³Reichenbach’s theory seems to be more popular in syntax than in semantics lately (at least since Hornstein 1990).

Typically, UT will be the origin of the reference frame, and, in past or future tensed sentences RT and ET will be displaced with respect to UT. Still, most times ET and RT will be in the relation of central coincidence (as in ‘I will be here tomorrow’). With more complicated combinations of tense and aspect, however, ET and RT do not have to coincide (and this is precisely what illustrates the need of RT). Consider the following examples, one in the past perfect, the other in the future perfect:

- (5) John had left the school.
- (6) John will have left the school.

In (5), ET occurs before UT, but also before a certain time that is different from UT: the utterance tells us that John’s leaving the school took place before the time we are talking about, which is also in the past relative to UT (RT, in this case, could be the time we got to the school – to facilitate this understanding, prefix (5) with ‘by the time we got there’). On the other hand, in (6), ET occurs after UT, but before a certain time that, again, is different from UT: the utterance tells us that John’s leaving the school took place before the time we are talking about, which is in the future relative to UT. The reference frames corresponding to (5) and (6) can be represented graphically by (5’) and (6’), respectively:

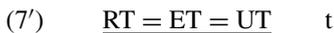


Examples (5) and (6) show that RT and ET can differ, but also that RT and UT need not coincide: in (5’) RT is situated before UT in the reference frame, while in (6’) RT is situated after UT. They can, however, coincide, as (4’), the reference frame corresponding to (4), illustrates.

Note that the fact that RT and UT need not, and do not often, coincide gives us already the means to explain cases (essentially similar to answering machine scenarios), in which a sentence in the simple present tense is used in a non-standard way. In most cases, when someone utters

- (7) I am here,

what is conveyed is that the speaker is at the location of the context *at the present time*. The reference frame corresponding to such an interpretation is represented as



However, (7) can be used non-standardsly. Recall the case of Anne and her colleague who comes in only on Thursdays at a certain time. Imagine that Anne wants to let her colleague know that she is in her office on Thursday when the colleague comes in. Trying to reduce her workload, Anne is leaving a post-it note with (7) on her colleague’s desk on Monday. The reference frame corresponding to such an interpretation of (7) is represented as



Here, RT and ET are located in the future with respect to UT, but the present tense is moved to the future (so to speak) to coincide with RT and ET, which are thus turned into the origin of the temporal reference frame. What is important to note is that the explanation we offer for this interpretation of (7) is facilitated by the framework adopted. Having three parameters, which stand for points in the time line, allows that the origin of the frame of reference can be moved to one or the other. However, our study of the behavior of ‘now’ will reveal that the origin of the frame of reference can only be moved to RT, that is, to the time we are talking about.

On the face of it, it may sound weird to dislocate UT from RT and ET in a present tense utterance, but the fact is that we can do it, and moreover, that we do it often enough. Usually, we use UT to fix the origin of the temporal frame of reference: time starts running from the moment we speak. So, RT and ET are located with respect to UT. But just as we may not use our own body as the center of a spatial frame of reference (more on this in Sect. 4), we may also not use UT as the center of a temporal frame of reference. Moving the temporal frame of reference forward or backwards is analogous to using an allocentric spatial frame of reference (centered not on our own body but on an external body).

On this construal, the alleged shifts in the use of temporal indexicals in present tense utterances are not motivated by non-standard uses of them in any case. What is going on, as we will immediately see, is motivated by where we place the origin of the temporal frame of reference with respect to RT. In this regard, Reichenbach (1947) is wrong in claiming that tenses determine time with reference to UT: denotation of time by tense depends on the origin of the frame of reference, which does not have to be UT.

3 The Proposal

The Reichenbachian framework sketched above has been extended to the treatment of temporal adverbials such as ‘at 3 pm’, ‘in 2 days’, ‘Monday’ etc. It is typical of such neo-Reichenbachian approaches to treat temporal adverbials, indexicals included, as modifiers of ET and/or RT. Temporal adverbials can modify the temporal reference of ET or of RT, and specify values for them (Hornstein 1990). To illustrate, consider

(8) Amina and Zara had left the house at 3 pm.

(8) can get two readings: according to one, Amina and Zara’s leaving took place at 3 pm; according to the other, Amina and Zara’s leaving took place before 3 pm. In the first interpretation “at 3 pm” modifies ET, while in the second it modifies RT (Demirdache and Uribe-Etxebarria 2007).

Importantly for us, according to this application of the Reichenbachian framework to temporal adverbials, indexicals differ from other adverbials in that they cannot modify the temporal reference of ET; rather, they seem to be always modifiers of the temporal reference of RT. To see this, consider

(9) Amina and Zara had left the house now.

(9) can only mean that Amina and Zara had already left the house before a certain time, which is interpreted as RT.⁴

Hornstein (1990) holds that ‘now’ is a temporal adverb that gives the value UT to RT. However, he also recognizes that other temporal moments can behave as temporal anchors. This idea could be put in terms of origins of reference frames: ‘now’ gives the value “origin of the reference frame” to RT. The origin of the reference frame is usually UT, and in such cases ‘now’ is effectively saying that RT coincides with UT. But the origin of the temporal frame of reference can be some other temporal anchor: when RT is displaced with respect to UT, ‘now’ marks RT as the origin of the frame of reference. A perhaps a more interesting way to put things is this: usually, a use of ‘now’ gives the hearer information about RT, saying that RT coincides with UT (“the time I’m talking about is the time of my speaking”). But in other cases the information that a use of ‘now’ gives presupposes an agreement about RT and has the effect of moving the origin of the temporal reference frame to RT (“let’s use the time we are talking about as point zero in the temporal line”). ‘Now’ marks a certain time as the zero point in the time line. As it always signals or picks up RT, this shows that only RT, and not ET, can be the origin of the reference frame.

With this in mind, we can explain the non-standard use of ‘now’ by answering machines, post-it notes etc. Thus, the reference frame of the interpretation of (1) played by an answering machine is the following:

(1') UT ET = RT t,

with ‘now’ picking up RT, which is the same as ET, both situated after UT and marking the origin of the reference frame (where we start counting from). By the same token, we can account for the historical present. The reference frame for the interpretation of (3) is represented as follows:

(3') ET = RT UT t

⁴See also: ‘He decided that he could not – definitely could not – tell Cousin Eunice any more about the white thing. It had been too disgusting. Anyway – whatever it was – it had left the house now.’ (Joan Aiken, *The Shadow Guests*). Indexicals other than ‘now’ seem to be able to modify both ET and RT simultaneously: for example, ‘She had left today/this week/this month’. But the case of ‘now’ shows they modify RT and not ET. The appearance that they also modify ET is due to the time stretch covered by ‘today/this week/this month’ etc., which can include both RT and a previous ET. Another indication that temporal indexicals go with RT and not with ET is shown by minimal pairs such as:

Yesterday, John had left a week ago.

*A week ago, John had left yesterday (Hornstein 1990, 26),

where it is shown that ‘yesterday’ is unable to modify the ET, and that indexicals, in contrast with anaphoric abverbials (‘a week ago’) do not seem to go well with ETs.

With ‘now’ picking up RT, which is again the same as ET, but this time both being situated before UT.

Interestingly, on this model we can also explain non-standard uses of ‘today’. In most cases, when someone utters

(10) I am here today,

what is conveyed is that the speaker is at the location of the context *on the current day*. The reference frame corresponding to such an interpretation is represented as

(10') $\underline{UT} \subseteq RT = ET \quad t$,

where UT, ET and RT refer now to intervals rather than moments of time (‘ \subseteq ’ signifying the *part of* relation). However, (10) can be used non-standardsly. Recall once more the case of Anne and her colleague. Imagine that Anne wants to let her colleague know that she is in her office on Thursday, which is the only day the colleague comes in. Trying to reduce her workload, Anne is leaving a post-it note with (10) on her colleague’s desk on Monday. The reference frame corresponding to such an interpretation of (10) is represented as

(10'')¹ $\underline{UT} \quad RT = ET \quad t$

Here, the intervals picked up by RT and ET crucially don’t overlap with the interval picked up by UT. Again, what is important to note is that the explanation available for this interpretation of (10), as well as for the answering machine cases and historical past, flows directly from the framework adopted (specifically, from the interplay between RT and UT and the origin of the frame of reference), and is thus not postulated to account solely for non-standard uses of ‘now’ or ‘today’. The only thing that is special in these cases is that the present tense is “moved” to where RT is – that is, we locate the origin of the temporal frame of reference in RT.⁵

Using a slightly modified Reichenbachian model, then, we could say that ‘now’, being always attached to RT, can do two things (but only two things). In standard uses, which exploit a standard temporal reference frame, with UT in its origin, ‘now’ provides information about where RT is located (namely in UT). In non-standard uses, however, ‘now’ provides information not about RT but about the origin of the temporal reference frame. It picks up RT (as it always does), but the information conveyed is that the origin of the temporal reference frame is where RT is (instead of conveying information about where RT is located). That is, by doing what it always does, i.e., picking up RT, ‘now’ can convey different kinds of information: about

⁵Hornstein (1990) makes a similar point: “These temporal anchors must be salient in the speech situation, which is why the speech time is so ready an anchor. However, other temporal moments can be exploited under the right circumstances, as is indicated by the use of tenses in historical narratives” (11). Yet, after noting this, he goes on to defend the view that temporal indexicals are always anchored to UT, such that they always “say” that RT coincides with UT. But this is not the case: apart from the answering machine cases, examples such as ‘Now they had left, Napoleon now faced his most crucial battle’, etc., show that ‘now’ does not give the value UT to RT.

where RT is, and, when RT is mutually known to speaker and hearer, about where the zero point in the temporal line is. The felicity of non-standard uses crucially depends on shared knowledge or assumptions about RT, as it is this kind of knowledge or assumptions that is exploited: “we both know, or we both assume, that I, the speaker, am talking about this time or interval. I invite you to consider such time as point zero in the time line”.

The account only needs to appeal to two Reichenbachian parameters, UT and RT, plus the idea that speaking about time always involves picking up a certain reference frame whose origin is placed somewhere in the temporal line. More often than not, authors simply assume that the origin of the reference frame is UT, but, as we have just seen, it need not be. However, this assumption has the effect of obscuring the fact that expressing time in language requires a point zero (an origin of the reference frame). On the other hand, speakers cannot pick up just any point in the time line as point zero and refer to it as ‘now’: for communication to succeed, ‘now’ can only refer to UT (or, as Horstein puts it, give the value UT to RT) or be anchored directly to RT itself when RT is mutually known to or assumed by speaker and hearer.

In the answering machine scenario, RT is the time when the message is played. The speaker knows that the hearer will know that she, the speaker, is talking about the time the hearer will listen to her message, so RT will be mutually known as soon as the message begins to be played. In this scenario, ‘now’, as well as the present tense, marks RT as the origin of the temporal reference frame. In the case of historical present, narrative clues make clear when RT is located. The choice of the present tense and the non-standard use of ‘now’ has the effect of situating point zero of the temporal line in RT. The case of the post-it note is not so straightforward. Anne’s colleague, who sees the post-it note on her table, maybe is not aware that the time Anne is talking about is precisely the time she is reading the note. However, her best guess is to assume that Anne is referring to that time, and so that RT is the moment when she reads the note; otherwise, the message does not make sense.⁶

4 ‘Now’ and Past Tense Constructions

Note that while this could eventually count as a satisfying explanation of what we do when we use ‘now’ non-standardly, it cannot be the end of the story about ‘now’. What has been said so far is an account of ‘now’ + *present tense* constructions. It

⁶In this, the view proposed resembles both “the coordination account” of demonstratives put forward in King (2014) (despite the fact that King explicitly excludes ‘now’ from the expressions his view applies to) and – perhaps more so, since ‘now’ is among the expressions considered – the “mutually-accepted perspective” view defended in Mount (2008, 2015). While both authors claim that coordination (in King’s case) and mutual acceptance of perspectives (in Mount’s case) are necessary for demonstratives and, respectively, indexicals to refer to a certain object, we want to remain neutral between such a claim and the weaker one that they are necessary for successful communication.

is this kind of constructions which have occupied philosophers, and for a reason: they are the constructions that question Kaplan's original account more directly. However, 'now' can appear in different linguistic environments. In (11) it appears with a past perfect, and in (12) with an imperfect:

(11) Amina and Zara had left the house now.

(12) Napoleon now faced his most crucial battle.

As far as we know, there is no account in the kind of literature that we are considering that deals with cases like these. However, they apparently show very clearly that 'now' cannot be an indexical that picks up UT. By the same token, they also show that 'now' cannot be an indexical that picks up tokening time, hearing time, or any other proposed parameter. Actually, if examples like (11) and (12) are taken at face value, they seem to refute most, if not all, extant accounts concerning 'now'. So, either they are explained away or they are taken as data that supports the idea that 'now' does not signal UT but RT. In (11) RT is in the past with respect to UT, but is located after the event of leaving occurred. The three parameters occupy different points in the time line, and it seems that 'now' goes with RT. What can we say about these, admittedly, infrequent, uses of 'now'? To answer this question let us talk a bit more about frames of reference, drawing on an analogy with locating ourselves in space.

It is a truism that for the interpretation of locational expressions such as 'on the left' we need the provision of a location. The default location is usually that of the speaker, but it need not be. Compare:

(13) The ball is on my left.

(14) The ball is on the left.

(13) is explicit about the origin of the spatial coordinates: it is the speaker. (14), however, by itself, is unspecific.

We can locate objects in space using different reference frames: egocentric, allocentric (intrinsic or relative), or absolute. The egocentric frame of reference locates things with respect to the speaker's body, an absolute frame of reference locates things with respect to certain environmental invariants and an allocentric frame of reference locates things with respect to some element in the environment which can either have a natural front, back, left and right (such as a person) or not (such as a tree). In the first case, the allocentric frame of reference is intrinsic, and in the second case it is relative. Thus, if someone utters 'The ball is on the left of the tree', she is using an allocentric frame of reference (the tree) but this frame of reference is non-intrinsic: the tree does not have a left or a right. Rather, she locates the ball on the left of the tree with respect to a certain viewpoint.

Locating times and events in the time line is similar to locating points and objects in space.⁷ We can treat UT-anchored utterances as similar to ego-based ones, except that using UT as a temporal anchor is much more frequent than using the body as a spatial anchor. But the UT anchor is just a default anchor. In face-to-face interactions, speaker and hearer will share a natural temporal frame of reference. This is not the case with respect to space: speaker and hearer do not usually share a natural/default spatial frame of reference. This is what can make (14) infelicitous. In any case, the point is that what is regularly done with respect to space can also be done with respect to time, such that if speakers and hearers begin to occupy different regions in the time dimension (as in the answering machine scenario), they can also choose to adopt an “allocentric” frame of temporal reference. Thus, we can use the temporal equivalent of allocentric frames of reference (and of course, we can use absolute coordinates).

The variety of allocentric frames of reference that can be used in communication seems to be much more restricted in time than in space. The spatial layout offers many objects that can be used as origins of spatial coordinates. The time line is less generous. The only “visible” points are UT, RT and ET. Sometimes, like in perfective constructions (‘John had left by the time we got there’), ET cannot be located with any exactitude. So, it seems that point zero has to be either UT or RT when RT does not coincide with UT. The usual counterexamples to the Kaplanian approach are cases in which the origin of the frame of reference is moved to RT: while in the case of (1) and (2) (answering machine and post-it notes cases) UT is in the past with respect to RT, in the case of (3) (historical present) UT is in the future.

These standard counterexamples would be the equivalent to using an intrinsic, non-egocentric, frame of reference. However, the point is that reference to time allows that non-intrinsic frames of reference can also be used. Thus, in (11) and (12) the origin of the frame of reference we use is located in the past-from-the-viewpoint-of-UT. Statements like these convey that the origin of the temporal frame of reference is situated in the past with respect to our time, just as when we locate the origin of a spatial frame of reference in an external body as seen from our perspective. In the latter case, things are ordered with respect to that object from our viewpoint; in the former, events are ordered with respect to a point in the past (RT) from the UT point of view. It is like, e.g., (12) is saying: now the action starts, this point in the past from our viewpoint is the starting point of what comes next.

Thus, also uses of ‘now’ that are embedded in past tense structures get an explanation using the Reichenbachian apparatus. In sum, uses of ‘now’ + *present tense* locate the origin of the reference frame at RT with no further ado (sometimes having the effect of “moving” the origin of the reference frame away from UT). Uses of ‘now’ + *past tense*, however, locate the origin of the frame of reference at RT as well, but treat RT as a point in time in the past with respect to UT, using thus a relative frame of temporal reference.

⁷This is not to say that we want to apply to ‘here’ the same treatment we applied to ‘now’. As we said, we are focusing here only on the latter, and we use the spatial case a supporting analogy.

5 Comparison with Other Views

The main claim of our proposal is that ‘now’ always refers to RT, although it can give two different kinds of information: that RT is at UT, or that point zero is at RT. RT, in turn, need coincide neither with UT nor with ET. We have shown how this flexibility helps with the problematic cases tackled in this paper.

But how different is this view from other views on the market? As we mentioned, our aim in this paper hasn’t been to engage with the multitude of solutions to the answering machine puzzle, but it would nevertheless be useful to show how it differs from at least some prominent views (or, more precisely, families of views). We also claimed that the present proposal both departs from and preserves some core claims of the Kaplanian orthodoxy. We start with spelling out the differences and the commonalities between our view and Kaplan’s.

One obvious way in which the two frameworks differ is that in our view ‘now’ always picks up RT, while the Kaplanian orthodoxy has it that ‘now’ always picks up the utterance time. But given that RT need not coincide with UT, even in present tense utterances, we can, while Kaplan cannot, hold that ‘now’ can refer to times other than the utterance time. Indeed, this has been the feature of the view that lies at the basis of our solution to the answering machine puzzle. Another way in which our view differs from Kaplan’s is that we allow some flexibility with respect to where we locate the zero point in the time-line. Whereas it is usual to locate it in UT, we might prefer to locate it somewhere else.

While we take these differences to be important, we would like to stress that the view proposed is similar to Kaplan’s in several, equally important respects. First, for us, as for Kaplan, ‘now’ has a *non-ambiguous character*. True, we need to assign a different character to ‘now’ than Kaplan does (e.g. ‘now’ picks up the reference time, and not the utterance time), but the difference can be considered minimal. Second, for us, as for Kaplan, the reference of ‘now’ is determined *in the context of utterance*. These two similarities between Kaplan’s view and ours are important in that they allow us to show that our view is less radical than other extant views on the market. For example, the fact that in our view ‘now’ has a non-ambiguous character differentiates and, to our mind, gives our view an advantage over views that postulate (certain types) of ambiguity. Thus, the view departs both from early claims that ‘now’ is ambiguous (e.g. Smith 1996), but also from more recent views according to which ‘now’ has a “shifty character” (Michaelson, 2013) – a claim that basically amounts to treating the word as polysemous. This is not so on our view.

The fact that the reference of ‘now’ is determined in the context of utterance also differentiates our view from a type of views that postulate additional contexts that are relevant in determining the reference of ‘now’. For example, Predelli (1998) claims that ‘now’ gets its reference from the “intended context”; Dodd and Sweeney (2010) and Cohen (2013) that it does so in the “context of interpretation” etc. In our view there is no need to postulate such context in addition to the context of utterance: ‘now’ gets its reference in the latter; it’s just that – given the flexibility of the framework – that reference need not coincide with UT. In a similar vein, the

view we propose also differs from the family of views that appeals to “deferred utterances”, “remote utterances” or other mechanisms for “making utterances at a distance” (Sidelle, 1991; O’Madagain 2014; Briciu 2018 etc.). In our opinion, and ignoring the slightly odd metaphysical character of such entities, there is no need for them: the original utterance is the one that should matter in determining the reference of ‘now’, and since the original utterance (the recording of the message, or the writing of the post-it note) happens in the context of utterance, that context is the only one we need.

One further issue that helps differentiate our view from others on the market is that of the role of intentions. What is the role played by speaker intentions in determining the reference of ‘now’ in this framework? As this question touches fundamental issues in semantics, we won’t pretend here to have a full-blown answer to it, but the quick and intuitive response is this: in our view, intentions play the role of establishing the zero-point in the temporal time line (basically, to use UT, as usual, or to use RT when RT does not coincide with UT).

However, is this answer enough to set out our view from others? For example, Perry (2006) has famously claimed that ‘now’ is not an automatic indexical because it can refer to the present moment, or to an age, or to an entire epoch etc. What the interval picked out by ‘now’ is essentially depends on the intentions of the speaker. The same, it could be said, happens on our view. But there’s a clear difference between our view and Perry’s: regardless of what interval ‘now’ picks out, in Perry’s view UT is *always* part of that interval. As should be clear by now, this is not true on our view: regardless of the length of the interval picked up by ‘now’, it is always RT⁸ – which need not (and in the relevant cases it *does* not) overlap with UT. In other words, our view does not rest content with merely claiming that ‘now’ is not an automatic indexical (a claim we adhere to), but to a more radical claim regarding the (non)coincidence of RT with UT.

Another view that puts significant weight on speaker’s intentions is Predelli’s (1998). According to Predelli, ‘now’ gets its reference at the *intended* context, which – as in our view – might differ from the utterance context in the relevant cases. What is then the difference between the two views?

The differences are quite substantial, we think. First, as mentioned earlier, on our view there is no need to postulate any additional context to than of utterance: ‘now’ gets its reference in the context of utterance. Second, while intentions do play a role, their role is more indirect, in that they determine the configuration of the reference frame, but not what parameter ‘now’ picks up. That is, it is possible to intend to use “non-standard” (or better, non-habitual) origins of reference frames, but ‘now’ will always be where RT is. In the moment she records her message, the speaker can choose between locating point zero at UT and speaking about RT in the future (‘I won’t be home’), or having point zero at RT (‘I’m not home’). The choice depends on what she believes is going to be more comprehensible. In the first case, she

⁸Here we envisage a version of the view that takes UT, ET and RT to be intervals, not moments of time.

intends to have point zero at UT; in the second, she intends to have point zero at RT. However, in the moment she uses ‘now’ in her message and tries to say something about RT in the future, she has to locate point zero in the future with respect to her utterance. Thus, the accusation of playing Humpty-Dumpty, leveled against Predelli, cannot be brought against our view. Finally, the view is committed to the claim that the reference of ‘now’ is determined by means of a linguistic rule – ‘now’'s character. This suggests that our view is more linguistically-driven than Predelli's,⁹ thus preserving one of Kaplan's main claims about ‘now’.

To finish comparing our view with other extant views, it would be instructive to see, briefly, how our view differs from that put forward by Mount (2008, 2015). Mount's view has at least two traits that we welcome (and which it shares with ours): it is very flexible – a crucial feature in solving the answering machine puzzle, and it offers a unitary explanation of both standard and non-standard uses of ‘now’. We have also made some remarks about ‘now’ getting a referent when some kind of coordination between speakers and hearers takes place – which is the main idea she advocates for. However, the advantage we take our view to have over Mount's is, again, its closeness to the Kaplanian orthodoxy. Mount (and, if Predelli (2011) is right, a great many of the proposed solutions to the answering machine puzzle) rejects “the propriety thesis” – an essential part of the Kaplanian framework, amounting to the idea that semantic evaluation of utterances should take place only in “proper contexts”, namely those in which the agent of the context is at the place and time of the context. Yet, we think that our view preserves the propriety thesis for, in our view, it remains true that the agent of the context is at the place and time of the context – that is, the time and place of the context *of utterance*, which is the only context we have and which is the one in which (in the case of the answering machine puzzle) (1) was recorded. Again, the fact that UT and RT don't coincide allow us to preserve the thesis while accounting for the puzzle.¹⁰

In sum, we take the advantages of the Reichenbachian view put forward here to be the following. First, our treatment of ‘now’ belongs to a systematic account of an entire class of expressions – namely, temporal adverbials. Second, the proposal makes use of a recognized and useful framework. This is compatible with saying that Reichenbach's original proposal might need to be amended in various respects; our proposal, as well as those of Hornstein, Demirdache and Uribe-Etxebarria and of other neo-Reichenbachians, need not be committed to the letter of Reichenbach's proposal. However, the important thing here is that, by applying the neo-Reichenbachian framework at stake, we are able to offer a unitary treatment both of standard uses of ‘now’ and of non-standard uses, such as that in answering machine scenarios, the historical present, past and future perfect tenses, and others.

⁹More recently, Predelli (2011) has argued that his view has room for linguistically-driven constraints as well.

¹⁰We also have an additional worry for all views postulating a single time parameter: how would complex tenses like the past or future perfect (illustrated by our examples (5) and (6) above) be accounted for? Such examples are easily handled in the Reichenbachian framework we adopt: indeed, it was devised precisely with this aim in mind.

Many other approaches lack this explanatory power. Third and finally, the view put forward departs from Kaplan, but at the end of the day only minimally, and is thus less radical than other (families of) views on the market, in that – for example – it doesn't postulate other types of contexts relevant for the interpretation of 'now', it doesn't appeal to deferred or other metaphysically dubious entities and it suitably limits appeal to intentions.¹¹ Of course, we haven't given any arguments that the alternative views mentioned are fundamentally mistaken; but to those who prefer the simplicity of the Kaplanian framework, adopting the Reichenbachian view put forward here looks like a very reasonable way to go.

Acknowledgments Agustin Vicente thanks Myriam Uribe-Etxebarria and Elena Castroviejo for comments and acknowledges the financial help of a grant from the Ikerbasque Foundation for Science. Dan Zeman thanks Adrian Briciu and Triinu Eesmaa for comments, as well as the audiences at the annual meeting of the *Canadian Association for Philosophy*, University of Quebec (4–7.06.2018) and *PhiLang06*, University of Łódź (12–14.05.2017). He acknowledges the financial help of the *Foundations and Methods of Natural Language Semantics* project (FFI2016-80636-P, University of Barcelona), of the *Lexical Meaning and Concepts* project (FFI2014-52196-P, University of the Basque Country) and of a "Lise Meitner" research grant (M2226-G24) from the FWF.

References

- Åkerman, J. (2009). A plea for pragmatics. *Synthese*, 170, 155–167.
- Briciu, A. (2018). Indexicals in remote utterances. *Philosophia*, 46, 39–55.
- Cohen, J. (2013). Indexicality and the puzzle of the answering machine. *The Journal of Philosophy*, CX, 5–32.
- Cohen, J., & Michaelson, E. (2013). Indexicality and the answering machine paradox. *Philosophy Compass*, 8, 580–592.
- Corazza, E., Fish, W., & Gorvett, J. (2002). Who is I? *Philosophical Studies*, 107, 1–21.
- Demirdache, H., & Uribe-Etxebarria, M. (2007). The syntax of time arguments. *Lingua*, 117, 330–366.
- Dodd, D., & Sweeney, P. (2010). Indexicals and utterance production. *Philosophical Studies*, 150, 331–348.
- Gauker, C. (2008). Zero tolerance for pragmatics. *Synthese*, 165, 359–371.

¹¹Let us also mention another advantage of our view that the other views don't seem to have: the possibility of accounting in a straightforward way for cases in which 'now' is bound. This can happen in two ways. First, 'now' interacts with adverbs of quantification such as 'usually'. In a sentence like 'Usually, John is here now', 'now' doesn't refer to the utterance time, but to a time that 'usually' quantifies over. We find it very plausible to hold that that time is RT (which in this particular example overlaps with ET). See also Predelli's (1996) main example. Second, binding can happen even in the absence of adverbs of quantification, given sufficient previous information. Mount (2008) gives the following example of a zookeeper who describes her daily activity: 'First I feed the giraffes. Then I make sure the elephant enclosure is free of debris. Now it's time to check on the gorillas...' (196). In the same vein, we take the bound time to be RT. We thank Peter Pagin for bringing this advantage of the view to our attention.

- Kaplan, D. (1989). Demonstratives: An essay on the semantics, logic, metaphysics, and epistemology of demonstratives and other indexicals. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan* (pp. 481–563). New York: Oxford University Press.
- Hornstein, N. (1990). *As time goes by: Tense and universal grammar*. Cambridge: MIT Press.
- King, J. (2014). Speaker intentions in context. *Nous*, 48, 219–237.
- Klein, W. (1995). A time relational analysis of Russian aspect. *Language*, 68, 525–552.
- Michaelson, E. (2013). Shifty characters. *Philosophical Studies*, 167, 519–540.
- Mount, A. (2008). The impurity of ‘pure’ Indexicals. *Philosophical Studies*, 138, 193–209.
- Mount, A. (2015). Character, impropriety, and success: A unified account of Indexicals. *Mind & Language*, 30, 1–21.
- O’Madagain, C. (2014). Indexicals and the metaphysics of semantic tokens: When shapes and sounds become utterances. *Thought*, 3, 71–19.
- Perry, J. (2006). Using indexicals. In M. Devitt & R. Hanley (Eds.), *The Blackwell guide to the philosophy of language* (pp. 314–334). Malden: Blackwell.
- Predelli, S. (1996). Never put off until tomorrow what you can do today. *Analysis*, 56, 85–91.
- Predelli, S. (1998). I am not here now. *Analysis*, 58, 107–115.
- Predelli, S. (2011). I am still not here now. *Erkenntnis*, 74, 289–303.
- Reichenbach, H. (1947). *Elements of symbolic logic*. New York: The Free Press.
- Romdenh-Romluc, K. (2006). I. *Philosophical Studies*, 128, 257–283.
- Ruffino, M. (2012). Deferred utterances and proper contexts. *Disputatio*, IV, 807–822.
- Sidelle, A. (1991). The answering machine paradox. *Canadian Journal of Philosophy*, 21, 525–539.
- Smith, Q. (1989). The multiple uses of indexicals. *Synthese*, 78, 167–191.

Distributed Utterances



Mark McCullagh

Abstract I propose an apparatus for handling intrasentential change in context. The standard approach has problems with sentences with multiple occurrences of the same demonstrative or indexical. My proposal involves the idea that contexts can be *complex*. Complex contexts are built out of (“simple”) Kaplanian contexts by ordered n -tupling. With these we can revise the clauses of Kaplan’s Logic of Demonstratives so that each part of a sentence is taken in a different component of a complex context.

I consider other applications of the framework: to *agentially* distributed utterances (ones made partly by one speaker and partly by another); to an account of scare-quoting; and to an account of a binding-like phenomenon that avoids what Kit Fine calls “the antinomy of the variable.”

Keywords Context-sensitivity · Indexicals · Scare-quoting · Binding

1 Introduction

In this talk I’m going to work with the idea that contexts can be *complex*. The contrast is with contexts as they occur in David Kaplan’s Logic of Demonstratives (1977). I’ll call those *simple* contexts, and the complex ones are what you get from them by means of ordered n -tupling. That is, the set of complex contexts is the closure, under ordered n -tupling, of the set of simple (Kaplanian) contexts.

What use are complex contexts? We need these things, I’ll argue, in order to do semantics for what I’ll call **distributed utterances**: utterances that are more than usually spread out in time, location or other contextual features, in such a way that makes a difference to their proper interpretation.

M. McCullagh (✉)

Department of Philosophy, University of Guelph, Guelph, ON, Canada
e-mail: mmculla@uoguelph.ca

Interpreting distributed utterances is important not only in its own right but because we have a common device by which we indicate that a non-distributed utterance is to be interpreted *as if it were* distributed: scare-quoting. Thus, an account of distributed utterances is useful as part of an account of scare-quoting, which is a very common device in contemporary writing.

First I'll explain why we need to work with complex contexts. Then, I'll explain how the Logic of Demonstratives can be modified in order to handle complex contexts. Finally I'll talk about applications. A treatment of scare-quoting is one, but there's also a sort of *binding* that is possible in the complex-contexts setting.

2 The Problem

We are familiar with the idea that (in Kaplan's terminology) the *content* of an utterance depends on its *context*. In this framework, an utterance of "I am here today" by me, in Warsaw, on June 16, 2016 has its content in virtue of those just-listed features of it. Following Kaplan we can distinguish between all the innumerable and unsummarizable features of concrete actual or possible utterances, which we might call the *worldly* contexts of those utterances, and the formal summaries of those features that we use in a semantic theory, which we might call *formal* contexts. Formal contexts are used as *representations* of the (relatively few) features of worldly contexts that are relevant to the proper assignment of contents to utterances made in those contexts.

All this is standard. Also standard is that by "utterances" we mean, utterances *of sentences*. Correspondingly, in Kaplan's formal system, no clause specifies more than one context, even for complex expressions.¹

But it's long been recognized that assigning a single context to an entire sentence is problematic in some cases.² Consider:

Two Heres The speaker utters the sentence "It's cold here, but it's warm here" while moving, such that the first utterance of "here" occurs in a chilly spot by the window and the second utterance of "here" occurs in a warm spot by the fireplace.

Two Nows The speaker utters the sentence "Now I'm starting my sentence, and now I'm not starting it but finishing it."

Five Yous ("Afterthoughts" p. 586) The speaker utters the sentence "You, you, you, and you can leave, but you stay," pointing to a different person with each utterance of "you."

¹As is well known (e.g. Rabern 2013, 402), a clause describing how a "monstrous" operator works does so; but we are not discussing the possibility of monsters in this talk.

²Kaplan himself noted this. In "Demonstratives" he discusses the sentence "That is that" (§§IX, XVII) and in "Afterthoughts" the sentence "You, you you and you can leave, but you stay" and sentences containing multiple occurrences of "today" (1989, 586–87).

Intuitively, these utterances are true. But Kaplan’s framework prohibits this result. For in that framework we assign contexts to entire sentences, meaning that there will be only *one* location determining a content for “here,” one day determining a content for “today,” one *demonstratum*, and so on.

So, taking a whole sentence in one context doesn’t always work. What to do?

2.1 *Braun’s Proposal*

One proposal has been made by David Braun (1996). Braun focuses on occurrences of multiple *demonstratives* in a sentence. His consideration of this phenomenon shows, he argues, that we need to change our conception of linguistic meaning. A demonstrative word on its own does not have a character, Braun argues. Rather its linguistic meaning—what speakers associate with it in virtue of understanding such a word—is such that it acquires a character only in conjunction with a demonstration. Braun summarizes his proposal by writing:

we get the following (rough) picture of the linguistic meaning of ‘that’.

linguistic meaning + demonstration \Rightarrow character

This feeds into the Kaplanian diagram:

character + context \Rightarrow content

So demonstratives have three sorts of meanings: linguistic meaning, character, and content. (156)

Braun’s is a sophisticated and multifaceted proposal which I do not have time in this talk to treat in depth. What I wish to note for now is that the argument for such a major change in our conception of linguistic meaning is weakened if there is a less revisionary way to address the problem on which that argument is based. I will argue that there is such a way.

3 **Revising the Logic of Demonstratives (LD)**

The basic idea behind the revision is that we expand the *sorts of contexts* the definition of truth in LD quantifies over. Where we are evaluating a complex *expression*, we use a complex *context*. Simple contexts are used only for simple expressions.

Here are some clauses from that definition, along with the modifications for complex contexts. I’m presenting just a few, in order to illustrate the idea. A full

revision of the system proceeds on the same lines.³ The unstarred numbers are for clauses in Kaplan's system, the starred counterparts are the revisions of them.

Conjunction

$$\models_{cw} \phi \wedge \psi \quad \text{iff :} \quad \models_{cw} \phi \quad \text{and} \quad \models_{cw} \psi \quad (4i)$$

$$\models_{\langle c_1, c_2 \rangle w} \phi \wedge \psi \quad \text{iff :} \quad \models_{c_1 w} \phi \quad \text{and} \quad \models_{c_2 w} \psi \quad (4i^*)$$

Atomic formulas

$$\models_{cw} \pi \alpha_1 \dots \alpha_n \quad \text{iff :} \quad \langle |\alpha_1|_{cw} \dots |\alpha_n|_{cw} \rangle \in \mathcal{I}_\pi(w) \quad (2)$$

$$\models_{\langle c_1 \dots c_n \rangle w} \pi \alpha_1 \dots \alpha_n \quad \text{iff :} \quad \langle |\alpha_1|_{c_1 w} \dots |\alpha_n|_{c_n w} \rangle \in \mathcal{I}_\pi(w) \quad (2^*)$$

The indexical "I"

$$|I|_{cw} = c_A \quad (12)$$

$$\text{If } c \in C, \text{ then } |I|_{cw} = c_A \quad (12^*)$$

The idea is that when evaluating a context-sensitive expression at a context-point pair, we do so by assigning separate contexts to its parts. This makes it possible for a conjunction, for example to be "taken in" not one context for the whole sentence, but one context for the first conjunct and one context for the second. (This means no context-sensitivity for the connector; we could add that if we liked, odd though it is.)

We should want to preserve in the new system all the truths of Kaplan's system. We can do this by adding some principles to our revision. (Recall that C is the set of original, Kaplanian contexts, and C^* is the closure of that set under ordered n -tupling: the set of *complex* contexts.)

$$\begin{aligned} \text{Equivalence :} \quad \forall c \in C^* : c &\sim \langle c, c \rangle \\ &\sim \langle c, c, c \rangle \\ &\sim \langle c, c, c, c \rangle \\ &\text{etc.} \end{aligned}$$

$$\text{Collapse :} \quad \forall c_1, c_2 \in C^*, w : \text{if } c_1 \sim c_2, \text{ then :} \quad \models_{c_1 w} \alpha \quad \text{iff} \quad \models_{c_2 w} \alpha$$

³For simplicity's sake I'm omitting reference to variable assignments, collapsing Kaplan's two points of evaluation—world and time—into one point i , and dropping Kaplan's distinction between terms for places and terms for people. Fuller presentation of the revised system is a work in progress.

The reason for these principles is that we'd like to still be able to say the things we say in the original LD, i.e. things of the form $\models_{cw} \phi$ where c is simple and ϕ is complex. We can do so by taking whatever complex context ϕ 's syntactic structure demands (according to the revised, starred clauses), populating it entirely with c —getting a result like $\langle c, \langle c, c \rangle \rangle$, say—then using these principles to “collapse” that complex context to its equivalent simple context c . That is, the principles let us infer from

$$\models_{\langle c, \langle c, c \rangle \rangle w} \phi$$

to

$$\models_{cw} \phi$$

So the move to complex contexts doesn't deprive us of any of the statements we'd like to make in the original Logic of Demonstratives.

4 Implications and Applications

4.1 *Implication: Some Contexts Don't Work for Some Formulas*

One thing with the revised system is that it is not the case, as it is with LD, that *every* context is suited to the evaluation of *every* formula. For example, “I” can't be taken in $\langle c_1, c_2 \rangle$. If it so happens that $c_1 = c_2$ then our Equivalence and Collapse principles will let us infer a statement of its evaluation; but even they are no help if $c_1 \neq c_2$. So for each formula there will be infinitely many contexts in which it cannot be taken.

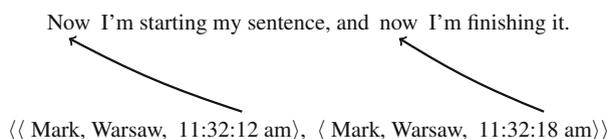
I think this is an unobjectionable consequence, since these are cases in which we shouldn't *want* the formula to be evaluable. Consider that example again. Suppose we want to insist that “I” be interpretable in the complex context $\langle c_1, c_2 \rangle$, where $c_1 \neq c_2$. What would our line be? Should it be the speaker of c_1 , or the speaker of c_2 , that is assigned as the value of “I” when taken in the complex context? There seems to be no principled reason to choose one over the other. Other ideas, like taking some sort of mereological sum of the two speakers, seem hopeless, if for no reason other than that the mereological sum of two speakers isn't the sort of thing we intuitively take possible referents of “I” to be.

There is a weaker principle, though, that is intuitively plausible and is preserved in the revised system. The principle is that every utterance is evaluable when taken in *the context of its utterance*. I think that the intuitive appeal of the stronger principle attaches entirely to the weaker principle too; there is no additional appeal that the stronger principle has. The idea of what something *would* mean, had it been

uttered in some other context, is not central to our everyday interpretative practice, even if careful philosophy of language gives us a vocabulary with which to talk about it.⁴

4.2 Application 1: Temporally/Spatially Distributed Utterances

We can handle Two Heres, Two Nows and Five Yous easily, while sticking with the classic Kaplanian idea that each *simple* context has one *demonstratum*, one time and one location. The diagram shows how this goes for Two Nows, exploiting the clause (4i*) I showed earlier; the other cases are handled on the same lines.



4.3 Application 2: Agentially Distributed Utterances

It is no quirk of Kaplan's formal system that the *agent* (speaker) of a context is simply one component of a context, formally on a par with its *time* and its *location*; formally, they *are* on a par. As concerns concrete speech acts however, they seem not to be. While we are ready to acknowledge that the time or location can change while an utterance is being produced, we are less ready to acknowledge that the speaker can change. Are complex contexts needed to handle what we might call *agentially* distributed utterances, the speaker of which changes part-way through their production?

Yes, for two reasons. One is that linguists have recognized that some utterances *are* agentially distributed. Ruth Kempson and co-workers (2004) have investigated the semantics of “shared utterances” within their “Dynamic Syntax” framework (2001). An attested example they give, from the British National Corpus:

- Daniel:** Why don't you stop mumbling and
Marc: Speak proper like?
Daniel: speak proper?

⁴Indeed, Kaplan famously deprecated—as “monsters”—operators that take as arguments not contents but ranges of contents *in other contexts* (1977, §VII). There is a lively debate about such operators, both over their formal legitimacy and over whether natural languages contain them. (See Schlenker (2003) and the many works citing it, and, recently, Rabern (2013).)

difference to the interpretation of indexical words in the enclosed portion. (The first example is from a recent biography of Stalin; the second is from the *New York Times*.)

After Denikin triumphantly entered Tsaritsyn [during Russia's Civil War] and attended services in its Orthodox cathedral, on July 3, he "ordered our armed forces to advance on Moscow." (Kotkin 2014, 326)

After a run through of "ideas I strongly reject," Bush finally got around to announcing that he was going to "talk about what we're for." (Gail Collins, *New York Times*, March 15, 2008)

In the first example the writer is indicating that the marked phrase is taken from some source in the historical literature, in which the words are presented as if uttered by some member of Denikin's forces in the Civil War. Thus the indication to the reader is to use such a person's context when interpreting the marked fragment. This contextual shift gives the right result for the occurrence of "our," despite the *overall* sentence being taken in a context that would assign to that word some plurality that includes the author of the book containing the example. Similarly in the second example. The writer is indicating that the quotation-marked words were uttered by President Bush; the contextual shift thereby suggested to the reader is one that makes sense of the indexicals, which otherwise would refer to the *New York Times* writer rather than to President Bush.

These examples show that taking scare-quotations as they're intended to be taken often involves the interpreter's changing, mid-way through the utterance, the context in which it is being taken. The complex contexts framework gives us a way to describe this. We thus handle scare-quoting *presemantically*—in our choice of context for the indicated part.

This is very different from handling scare-quoting *semantically*, as at bottom involving quotational reference to words. One representative semantic approach is that of Bart Geurts and Emar Maier. (They call the phenomenon in question "mixed quoting." There is no agreed-upon terminology for the cases in question. But evidently the approach would apply to the two examples I gave.) Here is their introductory description of their approach⁶:

(2) contains an instance of 'mixed' quotation. This is the species we shall be concerned with in the following pages:

(2) George says Tony is his 'bestest friend'.

In the case of mixed quotation, the meaning shift is, roughly speaking, *from α to 'what x calls ' α '*, where the value of x is determined by the context. For example, with $x = \text{George}$, the quotation in (2) is synonymous with the bracketed portion of (3):

(3) George says Tony is his [what George calls 'bestest friend'].

(Geurts and Maier 2003, 110)

⁶There are similar accounts in Brandom (1994), Benbaji (2004), and Recanati (2000).

The “meaning shift” makes theirs a *semantic* treatment. They assign a complex new meaning to a scare-quotation expression (that is, to the quotation marks plus the words they flank). That meaning involves *reference to* those words and *reference to* some speaker.

I haven’t time here to go into the relative strengths and weaknesses of pre-semantic and semantic treatments. But one large problem for the semantic approaches, I think, is that by building in hidden references to words and speakers, such treatments entail that these references should be available for anaphoric uptake later in the sentence, or in a following sentence. A look at many examples of scare-quoting, however, shows that attempts at such anaphoric uptake are infelicitous.⁷

So, one other use for the complex contexts framework is in an account of a *somewhat* quotation-like device that has proved challenging to account for under the assumption that each sentence is to be taken in one, simple Kaplan context. Given that mid-sentence shift in context is *possible*—which our prosaic examples at the beginning illustrate—it should not surprise us at all that for a variety of reasons, writers and speakers exploit our interpretative facility with it. They do so in their production of sentences that demand interpretation *as if* there were such a shift.

4.5 Application 4: “Co-indexing” by Context

One of Kit Fine’s arguments for semantic relationism starts from consideration of “an antinomy concerning the role of variables” (Fine 2003, 605) in formal logic. While the variable letters “*x*” and “*y*” seem to have the same role (in some pretheoretic sense of “semantic role”) in some formulas, it seems they play very different roles in others. In “ $x > 0$ ” and “ $y > 0$,” for example, they seem to have the same role. But in “ $x > y$,” on the other hand, they seem to have different roles (otherwise that expression would itself have the same role as “ $x > x$ ”). It is not obvious how to reconcile these two claims.

The antinomy relies on the assumption that the language employs more than one variable letter. Of course, this is a standard assumption, stated in every logic textbook:

First-order logic assumes an infinite list of variables so that we will never run out of them, no matter how complex a sentence may get. (Barker-Plummer et al. 2011, 231)

The reason for all the distinct variables is so that a quantifier prefixed to an open formula can bind some, but not all, of its free variables. Otherwise one quantifier would bind all free variables in a formula, and there would be no possibility of multiple (non-vacuous) quantification.

⁷I make this point more fully in McCullagh (2017).

Against this background it is notable that the complex contexts framework makes possible a form of binding—something, that is, that systematically reduces the assignment-dependency of semantic values—that allows for multiple quantification without involving multiple variable letters. The basic idea is to make *differences in context* do the work of distinctness of variable letters. *Occurring in the same (simple) context* is the relation that relates a particular quantifier occurrence (in a formula) to a particular occurrence of a free variable. This is something we can do only once we can take different parts of a formula in different contexts. I will briefly sketch the idea, presenting it as a variation on the familiar Tarskian account.

4.5.1 The Tarskian Account of Binding

On the Tarskian treatment, the truth value of a quantified formula, relative to one variable assignment, depends on the truth values of the formula with the initial quantifier stripped off, relative to *variant assignments*:

$$\models_{cfw} \forall x_n Fx_n \text{ iff : } \left\{ \begin{array}{l} \models_{cf_a^n w} Fx_n \\ \models_{cf_b^n w} Fx_n \\ \models_{cf_c^n w} Fx_n \\ \models_{cf_d^n w} Fx_n \\ \vdots \end{array} \right\} \text{ for all } a, b, c, d \dots \text{ in the domain}$$

Now suppose...

- ... we wanted to think of a variable assignment as something done by context. (This is somewhat plausible for occurrences of free variables.)
- ... and we wanted to work with only one variable rather than with infinitely many.⁸

If we're thinking of variable assignments as done by context, this amounts to having different occurrences of the same variable letter *be taken in different contexts*. We can have one context be used simply to “attach” to whatever quantifier we are using, and also to “attach” to whatever *occurrences* of variables that we wish that quantifier to bind. Like so:

4.5.2 “Binding” by Context

$$\models_{(c_1, c_2)w} \forall Fx \text{ iff:}$$

⁸The same points will apply if we want to work with some small number of variables rather than infinitely many—more like natural language.

$$\left. \begin{array}{l} \models_{c_2(c_1^a)w} Fx \\ \models_{c_2(c_1^b)w} Fx \\ \models_{c_2(c_1^c)w} Fx \\ \models_{c_2(c_1^d)w} Fx \\ \vdots \end{array} \right\} \text{for all } a, b, c, d \dots \text{ in the domain}$$

- c_1 is the *simple* context that “ \forall ” is being taken in.
- c_2 is the (possibly complex) context that “ Fx ” is being taken in.
- $c_2(c_1^a)$ is the context just like c_2 except that wherever c_1 is a component of it, it is replaced by c_1^a , which differs from c_1 in that our—one!—variable, x , is assigned a .

This gives us the effect of variable binding but with only one variable letter.

This is just the briefest of sketches. All I mean to have done here is to justify the claim that it is *possible* for there to be an account of variable binding on which the “antinomy of the variable” does not arise. Since part of the argument for semantic relationism is that it offers a solution to that antinomy, the “binding by context” approach should figure in an overall weighing of the plausibility of semantic relationism.

5 Conclusion

Working with complex contexts allows us to:

- handle the cases in which there is, intuitively, a mid-sentence shift in context (“distributed utterances”);
- tell a story about scare-quoting that handles indexicals in the scare-quoted material and doesn’t have the problems of semantic approaches;
- tell a story about variable binding, in which contexts do the work of subscripting, and the “antinomy of the variable” does not arise.

References

Almog, J., John, P., & Wettstein, H. (Eds.). (1989). *Themes from Kaplan*. Oxford: Oxford University Press.

Barker-Plummer, D., Barwise, J., & Etchemendy, J. (2011). *Language, proof, and logic* (2nd ed.). Stanford: Center for the Study of Language and Information.

Benbaji, Y. (2004). A demonstrative analysis of “open quotation”. *Mind and Language*, 19, 534–547.

- Brandom, R. (1994). *Making it explicit*. Cambridge, MA: Harvard University Press.
- Braun, D. (1996). Demonstratives and their linguistic meanings. *Noûs*, 30, 145–173.
- Fine, K. (2003). The role of variables. *Journal of Philosophy*, 100, 605–631.
- Geurts, B., & Maier, E. (2003). Quotation in context. *Belgian Journal of Linguistics*, 17, 109–128.
- Kaplan, D. (1977). Demonstratives. In Almog et al. (1989) (pp. 481–614).
- Kaplan, D. (1989). Afterthoughts. In Almog et al. (1989) (pp. 567–614).
- Kempson, R., Meyer-Viol, W., & Gabbay, D. (2001). *Dynamic syntax: The flow of language understanding*. Oxford: Blackwell.
- Kotkin, S. (2014). *Stalin: Volume I: Paradoxes of power, 1878–1928*. New York: Penguin.
- McCullagh, M. (2017). Scare-quoting and incorporation. In P. Saka & M. Johnson (Eds.), *The semantics and pragmatics of quotation* (pp. 3–34). Dordrecht: Springer.
- Purver, M., & Kempson, R. (2004). Incrementality, alignment and shared utterances. In J. Ginzburg & E. Vallduví (Eds.), *Catalog'04: Proceedings of the eighth workshop on the semantics and pragmatics of dialogue* (pp. 85–92). Barcelona: Department of Translation and Philology, Universitat Pompeu Fabra, Barcelona.
- Rabern, B. (2013). Monsters in Kaplan's logic of demonstratives. *Philosophical Studies*, 164, 393–404.
- Recanati, F. (2000). *Oratio obliqua, oratio recta: An essay on metarepresentation*. Cambridge, MA: MIT Press.
- Schlenker, P. (2003). A plea for monsters. *Linguistics and Philosophy*, 26, 29–120.

Demonstratives in First-Order Logic



Geoff Georgi

Abstract In an earlier defense of the view that the fundamental logical properties of logical truth and logical consequence obtain or fail to obtain only relative to contexts, I focused on a variation of Kaplan's own modal logic of indexicals. In this paper, I state a semantics and sketch a system of proof for a first-order logic of demonstratives, and sketch proofs of soundness and completeness. (I omit details for readability.) That these results obtain for the first-order logic of demonstratives shows that the significance of demonstratives for logic exceeds their behavior as rigid designators in counterfactual reasoning, or reasoning about alternative possibilities. Furthermore, the results in this paper help address one common objection to the view that logical truth and consequence obtain only relative to contexts. According to this objection, the view entails that logical consequence is not formal.

Keywords Demonstratives · First-order logic · Proof theory · Coordination · Kaplan

1 Introduction

David Kaplan's original investigations into the logic of indexicals offer valuable insights into the relationships between logic, meaning, and modality (Kaplan, 1979, 1989). Kaplan also tried to incorporate true demonstratives into his logic, via his artificial 'dthat', but David Braun (1996) and Nathan Salmon (2002) have argued that the syntax and semantics for 'dthat' fails as a theory of the demonstrative pronoun 'that' in English. Kaplan's other alternative is to consider adding infinitely many 'that's into a language, distinguishing each with a subscript: 'that₁', 'that₂', etc. This too fails as an account of the English demonstrative 'that', which is only

G. Georgi (✉)

Department of Philosophy, West Virginia University, Morgantown, WV, USA
e-mail: geoff.georgi@mail.wvu.edu

one word.¹ An adequate semantics for ‘that’ in Kaplan’s framework must allow distinct syntactic occurrences of the same word ‘that’ to differ in content relative to the same context. I sketch such a semantics in Sect. 3.

One result of Kaplan’s failure to incorporate true demonstratives satisfactorily into logic and semantics is that some of the deepest insights into logic his theory reveals have gone unrecognized. In this paper, I introduce a first-order logic for demonstratives in order to argue that the fundamental logical properties of logical truth and logical consequence for languages containing demonstratives obtain or fail to obtain only relative to contexts. In an earlier paper defending this thesis, I used a variation of Kaplan’s own modal logic of indexicals (Georgi, 2015). For the theorems in Sect. 5, the focus on first-order logic allows us to simplify proofs.

But the goal of this paper is not merely to present new results in the logic of demonstratives. The technical results in Sect. 5 play several supporting roles in the philosophical arguments that follow. Most importantly, that the theorems obtain for the first-order logic of demonstratives shows that the significance of demonstratives for logic exceeds their behavior as rigid designators in counterfactual reasoning or modal logic.²

The results in this paper also play a role in responding to one common objection to the thesis of this paper.³ The objection is that on the resulting theory, logical consequence is not formal in the fundamental sense in which logical consequence is supposed to be formal (whatever that fundamental sense is). This objection requires an assumption, some variation of which is common but usually left implicit in discussions of logical form, that I call in Sect. 4 the *coordination in context assumption*. This assumption raises the challenge that the logic of demonstratives presented here does not reflect anything deep or significant about logic. But once aired, and with some competitors considered, the coordination in context assumption is revealed to be somewhat bereft of justification.

To illustrate this claim, in Sect. 5 I sketch an argument for the relevant interpretation of the coordination in context assumption. This argument relies on a restricted understanding of soundness and completeness results for systems of first order logic. In response, I provide a system of proof for first-order languages containing demonstratives, and show that the resulting system is sound and complete with respect to the definition of logical consequence relative to a context in Sect. 3. These results undermine the argument for the coordination in context assumption,

¹The premise that there is only one demonstrative pronoun ‘that’ in English explicitly rules out theories according to which a use of ‘that’ to refer to Lone Pine Peak, and a use of ‘that’ to refer to something other than Lone Pine Peak, are uses of lexically distinct demonstratives (Gauker, 2014).

²The logical issues in this paper are also independent of the status of demonstratives as quantifiers (King, 2001, Hawthorne and Manley, 2012), individual concepts (Elbourne, 2008), or devices of direct reference (Kaplan, 1989, Braun, 2008, Georgi, 2012). Any such semantics must allow distinct occurrences of a true demonstrative to differ in content relative to the same context.

³Common in the sense that it, or something in the vicinity, is the objection I meet most frequently when presenting this material.

and further confirm the utility of understanding logical consequence, for languages containing demonstratives, as holding or not holding only relative to contexts.

2 Contexts and Logic

As a theoretical term, ‘context’ is used in several ways in contemporary linguistics and philosophy (Predelli, 2005, Stalnaker, 2014). In this paper, we are concerned with contexts as introduced into logic and semantics by Kaplan, who distinguishes between different roles played by elements of indices in earlier semantic and logical theories (Lewis, 1970, Montague, 1974). Some parameters of an index reflect features of the physical environment in which language is used. Others are required strictly for formal purposes, because semantics requires parameters that can be shifted. For an operator analysis of tense, for example, there must be a temporal parameter available to be shifted by tense operators (such as ‘it was the case that’). But the temporal features of a use of a tensed expression do not shift in virtue of the interpretation of the tense operator used (Lewis, 1981). The contexts of Kaplan’s formal theory represent features of the physical environments of utterances that are not shifted in semantics (but see Sect. 3). Kaplan called the parameters required by shifty operators *circumstances of evaluation*.⁴ Kaplan’s distinction between contexts and circumstances provided a robust philosophical foundation for earlier work both on two-dimensional modal logic (Seegerberg, 1973) and on double indexing (Kamp, 1971), and offered a still influential account of the relationships between logic, meaning, and modality.

In Kaplan’s theory, an indexical like ‘I’ or ‘now’ is assigned what he calls a *character*, which he identifies as the meaning of the indexical. This character determines for each context what Kaplan calls the *content* of the indexical relative to, or in, the context.⁵ This content, together with a circumstance of evaluation, determines the extension of the indexical at the context and circumstance. Thus the same indexical may differ not only in extension relative to different contexts and circumstances, but also in content relative to different contexts, while having the same meaning (character) throughout. In this way, Kaplan gives an intuitive semantic account of both a case in which Barack Obama and I utter (1), and a case in which I utter (1), and Obama utters (2) speaking only to me:

I am hungry. (1)

You are hungry. (2)

⁴It is now standard to use *index* for the collection of shifting parameters of a semantic theory.

⁵In Kaplan’s formal system, this content is an intension, but in his informal discussion the content of an expression in a context is the contribution the expression makes to the structured proposition expressed in that context by any sentence in which the expression occurs (Kaplan, 1989, pp. 494–496).

In the first case, Obama and I use the same sentence, with the same meaning, to say different things. In the second, we use different sentences, with different meanings, to say the same thing. Kaplan's theory is the foundation of a compelling account of these and a host of other basic linguistic facts.⁶

On the basis of the content of an expression relative to a context, Kaplan is able to give a systematic account not only of truth relative to a context for sentences containing indexical expressions, but also of logical truth for sentences containing indexicals. Assuming (as Kaplan does) that a context c uniquely determines a circumstance of evaluation, a sentence ϕ is true in a context c if and only if the extension of ϕ relative to c and the circumstance of evaluation of c is 1, or the True. (We might more naturally say that the content of ϕ relative to c is true at the circumstance of c .) Finally, a sentence ϕ of a language L is logically true if and only if ϕ is true in all contexts (in all models of L), and ϕ is a logical consequence of ψ_1, \dots, ψ_n if and only if ϕ is true in every context (of every model of L) in which all of ψ_1, \dots, ψ_n are true. This allows Kaplan to offer a persuasive account of the intuitive logical relations between (3), (4), and (5):

Beauty is sleeping now. (3)

It is not the case that Beauty is not sleeping now. (4)

Beauty is not sleeping now. (5)

(4) is a logical consequence of (3), because (4) is true in every context of every model in which (3) is true. Similarly, (3) and (5) are logically inconsistent, because there is no context in any model in which they are both true.

Logical truth is thus on Kaplan's view a property that a sentence like (6) in a language has or fails to have absolutely:

If I hike, then I hike. (6)

(6) is logically true because it comes out true in every context of every model of any relevant fragment of English. But this brief and intuitive explanation conceals an important philosophical point: the logical truth of (6) is guaranteed because the content of 'I' is fixed relative to a context, so in every context the truth of (6) reduces to a matter of propositional logic. Reflecting their origins in two-dimensional modal logic, contexts become merely one more parameter over which we generalize to obtain logical truth.

It is precisely this last feature of Kaplan's logic that I reject. I do so on the grounds that given this feature of Kaplan's logic, we cannot explain the logical behavior of

⁶Stalnaker developed very similar ideas around the same time (Stalnaker, 1972). Lewis disputes Kaplan's data (Lewis, 1981). Stalnaker offers to my mind a compelling response to Lewis (Stalnaker, 2014).

what Kaplan calls ‘true demonstratives’ (Kaplan, 1989, 490). True demonstratives are *referentially promiscuous*. Intuitively, an expression is referentially promiscuous if and only if as a matter of semantics it can be used more than once in the same sentence to refer to different things. More carefully, we may define referential promiscuity as follows:

REFERENTIAL PROMISCUITY

An expression *e* of a language L is referentially promiscuous if and only if there are distinct free occurrences O_1 and O_2 of *e* in a sentence *s*, and some context *c*, such that the content of O_1 relative to *c* (what O_1 contributes to the content of *s* relative to *c*) is distinct from the content of O_2 relative to *c*.⁷

For example, I might utter (7), while pointing first at a hiker, and then at someone who does not hike:

If that person hikes, then that person hikes. (7)

The referential promiscuity of demonstratives allows me to do this. It is a feature of the semantics of demonstratives that distinct occurrences are not semantically constrained in any way to refer to the same object or individual relative to the same context. Thus relative to the context of my use, the first occurrence of ‘that person’ in (7) refers to a hiker, while the second occurrence refers to someone who does not hike. Relative to this context, (7) is false. On Kaplan’s conception of logic, the existence of even one such context is sufficient to undermine the status of (7) as a logical truth. Thus (7) is not logically true.⁸

But there are also uses of (7) in which a speaker is guaranteed to say something true, such that this guarantee is as strong as if the speaker had uttered (8):

If Ranger Shelton hikes, then Ranger Shelton hikes. (8)

In other words, there are cases in which a speaker uses (7) in such a way that it seems logically true as used. Suppose, for example, that I utter (7) to A, pointing

⁷Two points about this definition: (i) My use of ‘referential promiscuity’ differs from Arthur Sullivan’s (Sullivan, 2013, ch. 4.4). He uses ‘referentially promiscuous’ to characterize context-sensitive expressions generally, whereas I reserve it for those expressions distinct occurrences of which can differ in content relative to the same context. I was unaware of Sullivan’s use when I first introduced my use of the term, and I have grown to fond of my use to change it. (ii) The restriction to free occurrences rules out consideration of bound variables. For discussion of the relationship between variables and demonstratives, see Georgi.

⁸Yagisawa argues that this kind of example arises even for the purest of indexicals like ‘I’, and for multiple occurrences of quantifiers (Yagisawa, 1993). His response to such examples seems to me far more radical than mine: Yagisawa accepts the consequence that any such sentence undermined by context-shifts is not logically true, and hence that all sentences and inferences so undermined should be expunged from a “pure” logic. Yet the resulting purification of logic strikes me both as an impoverishment of logic and as a misunderstanding of the significance of demonstratives. Refusing to consider logical consequence relative to a context, Yagisawa can only bar demonstratives from logic. I discuss Yagisawa’s view further in Georgi (2015).

the whole time at Shelton Johnson, who is standing nearby in clear sight, and that A is attentively following my gestures. Kaplan's treatment of logical truth as absolute, according to which (7) is not logically true, leaves us with no resources to explain the apparent logical truth of such uses.⁹

A similar observation holds for logical consequence. Suppose, for example, that we are standing in Lone Pine, California, facing the eastern escarpment of the Sierra Nevada. Pointing at Mt. Whitney, I utter (9):

Nothing is taller than that. (9)

Then pointing first at Lone Pine Peak, and then at Mt. Whitney again, I utter (10):

So, that is not taller than that. (10)

I have reasoned well. You can follow my argument and recognize its validity. Yet if demonstratives are referentially promiscuous, then it is a trivial exercise to construct a context relative to which (9) is true while (10) is false. Given even one such context, (10) is not a logical consequence of (9). As a result, we have no explanation of the apparent validity of the argument as I have used it.

Examples like these show that Kaplan's conception of logical truth and logical consequence is inadequate for a logic of true demonstratives. But rejecting Kaplan's conception of logical truth and consequence does not require abandoning the use of contexts in logic. Following Kaplan, I assume that sentences containing demonstratives are true only relative to contexts. In the semantics that follows, each context includes a sequence of parameters reflecting features of the use of demonstratives in conversation. This is precisely the role of contexts in Kaplan's *semantic* theory, though demonstratives require a broader understanding of those parameters than Kaplan develops. (I return to this point at the end of Sect. 3.) Furthermore, I agree with Kaplan that logical truth and logical consequence require generalizing over contexts. Where Kaplan's theory goes wrong is in limiting the role of contexts in logic. To repeat: I reject the view of contexts in logic as *merely* one more parameter over which we generalize to obtain logical truth.

3 A First-Order Logic of Demonstratives

In place of Kaplan's proposal, I suggest that a sentence like (7) is logically true, and that (10) is a logical consequence of (9), relative to some contexts but not relative to others. To be logically true relative to a context *c* requires that a sentence or formula

⁹Kaplan considers the skeptical possibility in which the speaker is ignorant of a "switcheroo" in which a powerful deceiver changes the scene mid-utterance. Such skeptical worries, if taken seriously, would undermine any study of logic. (How do we know that the meanings of our words aren't changed mid-argument by some omnipotent deceiver?) Thus I set such skeptical worries aside here.

be true in every context c' (in every model M) such that the same occurrences of demonstratives are linked in c' as are linked in c . These links between occurrences, or coordination schemes, are determined by context. In this section, I introduce a formal language containing demonstratives, provide a logic that develops this suggestion, and show how the logic applies to sentences like (7), and to arguments like (9), therefore (10).⁷

I characterize the first-order logic of demonstratives, or FOLD, semantically, delaying discussion of proof theory to Sect. 5. We suppose infinitely many variables v_1, v_2, \dots , infinitely many constants $\alpha_1, \alpha_2, \dots$, and for each n , sufficiently many n -place predicates P_1^n, P_2^n, \dots . We use the traditional '=' for the two-place identity predicate, which we interpret as a logical constant. In addition, we add the demonstrative ' δ '. Variables, constants, and ' δ ' are terms. We assume negation ' \neg ', the material conditional ' \supset ' and biconditional ' \equiv ', and universal quantification ' \forall ' as primitive.¹⁰ We use parentheses for punctuation. Formulas are formed in the usual way.

A model for FOLD is a triple $\langle C, D, I \rangle$, where D is a domain, I is an interpretation of the constants and predicates (that assigns the identity relation over D to '='), and C is a set of contexts, where each context is a pair $\langle d, R_c \rangle$ such that (i) d is a sequence of members of D , (ii) R_c is an equivalence relation over the positive integers, and (iii) for any positive integers i, j , if $R_c i j$, then $d_i = d_j$. After Fine (2003, 2007), I call R_c the *coordination scheme* for c . For any context c , the *coordinated variants* of c are those contexts c' such that $R_c = R_{c'}$.¹¹

Our semantics for ' δ ' has to allow distinct occurrences to differ in content relative to the same context. This requires that we have some means of tracking occurrences of ' δ '. The simplest way to accomplish this is to add an occurrence-tracking parameter to the definition of truth in a model.¹² Because our language is first-order, the result is that formulas of FOLD are true or false in a model relative to a context c , assignment f of values to variables, and occurrence-tracking parameter

¹⁰Even at this early stage, a significant difference arises between FOLD and standard languages of predicate logic. In FOLD, we have to be careful in our choice of primitive notation. In particular, we cannot get by with only one primitive connective (either ' \uparrow ' or ' \downarrow '). Assuming the standard practice of introducing other connectives by means of notational abbreviations, the result would affect the number of occurrences of a demonstrative in a formula. The demonstrative ' δ ' occurs twice in ' $(P\delta \downarrow P\delta)$ ', but only once in ' $\neg P\delta$ '. Given the semantics below, this can affect the interpretation of the formula relative to a context. To be an acceptable notational abbreviation the definiendum in the definition of a logical connective must contain precisely as many occurrences of ' δ ' as the definiens, and these occurrences must bear similar structural relations to one another in both the definiendum and definiens.

¹¹Paul Hovda first suggested that I incorporate coordination schemes directly into contexts at the Northwest Philosophy Conference in October 2012, though I now suspect that James Higginbotham had been trying to get me to see it years before.

¹²Pickel et al. (2018) suggest that context itself shifts, and they seem to have something like linguistic context in mind, in contrast to the extra-linguistic contexts of Kaplan's theory. I think this is largely correct, but for this paper it is easier to separate the occurrence-tracking parameter from context. The differences between these proposals are irrelevant to the purposes of this paper.

i , and terms of FOLD denote objects in a model relative to a context c , assignment f , and occurrence-tracking parameter i . Crucially, for each expression of FOLD we have to state not only the conditions under which it is true, or what it denotes, but also the effect that it has on the occurrence-tracking parameter. Formally, we introduce a function s from integers and expressions to integers. Thus each clause in our definitions of truth and denotation relative to a model requires two parts:

Definition 1 (Denotation in a model) Let ‘ $| t |_{cfi}^M$ ’ abbreviate ‘the denotation of t in the model M , relative to the context c , assignment f , and occurrence-tracking parameter i ’:

- D1 a. For any variable v , $| v |_{cfi}^M = f(v)$
 b. For any variable v , $s(v, i) = i$
 D2 a. For any (non-demonstrative) constant α , $| \alpha |_{cfi}^M = I(\alpha)$
 b. For any (non-demonstrative) constant α , $s(\alpha, i) = i$
 D3 a. $| \delta |_{cfi}^M = d_i$
 b. $s(\delta, i) = i + 1$

Definition 2 (Truth in a model) Let ‘ $\models_{cfi}^M \phi$ ’ abbreviate ‘ ϕ is true in the model M relative to the context c , assignment f , and occurrence-tracking parameter i ’:

- T1 a. $\models_{cfi}^M \ulcorner P^n t_1 \dots t_n \urcorner$ iff

$$\left(| t_1 |_{cfi}^M, | t_2 |_{cf s(t_1, i)}^M, \dots, | t_n |_{cf s(t_{n-1}, \dots, s(t_2, s(t_1, i)))}^M \right) \in I(P^n)^{13}$$

 b. $s(\ulcorner P^n t_1 \dots t_n \urcorner, i) = s(t_n, s(t_{n-1}, \dots, s(t_2, s(t_1, i))))$
 T2 a. $\models_{cfi}^M \ulcorner \neg \phi \urcorner$ iff $\not\models_{cfi}^M \phi$
 b. $s(\ulcorner \neg \phi \urcorner, i) = s(\phi, i)$
 T3 a. $\models_{cfi}^M \ulcorner (\phi \supset \psi) \urcorner$ iff either $\not\models_{cfi}^M \phi$ or $\models_{cf s(\phi, i)}^M \psi$
 b. $s(\ulcorner (\phi \supset \psi) \urcorner, i) = s(\psi, s(\phi, i))$
 T4 a. $\models_{cfi}^M \ulcorner (\phi \equiv \psi) \urcorner$ iff either
 i. $\models_{cfi}^M \phi$ and $\models_{cf s(\phi, i)}^M \psi$, or
 ii. $\not\models_{cfi}^M \phi$ and $\not\models_{cf s(\phi, i)}^M \psi$
 b. $s(\ulcorner (\phi \equiv \psi) \urcorner, i) = s(\psi, s(\phi, i))$
 T5 a. $\models_{cfi}^M \ulcorner \forall v \phi \urcorner$ iff for every $o \in D$, $\models_{cf_v^o}^M \phi$
 b. $s(\ulcorner \forall v \phi \urcorner, i) = s(\phi, i)$

In T5, f_v^o is the assignment function that differs from f at most in that it assigns o to v . Now let E be any equivalence relation over contexts. A formula ϕ is E -logically true relative to a context c if and only if ϕ is true in every model M relative

¹³Corner quotes in this paper are used for Quinean quasi-quotation (Quine, 1951, §6).

to every context of M that bears E to c , any assignment f , and the occurrence-tracking parameter 1. If E is any universal relation, such as *being either distinct from or identical to*, then a formula ϕ is E -valid relative to any context c if and only if it is valid in Kaplan’s logic of demonstratives LD. Other equivalence relations yield different logics.

Our use of demonstratives like ‘this’ and ‘that person’ in reasoning generates a natural equivalence relation between contexts. In reasoning, we intend certain uses of demonstratives to be linked in some way that guarantees coreference. The imagined use above of (7) while pointing continuously at Shelton Johnson offers an example:

If that person hikes, then that person hikes. (7)

The coordination schemes of contexts as defined above offer a formal model of this phenomenon we encounter in reasoning. Contexts that share the same coordination scheme will ensure that the same pairs of occurrences of demonstratives are coreferential. Above, I introduced the coordinated variants of a context: any two contexts are *coordinated variants* of each other if and only if they share the same coordination scheme. Being coordinated variants of one another is an equivalence relation over contexts. We may therefore define logical truth relative to a context as follows:

Definition 3 (Logical truth relative to a context¹⁴) A (closed) formula ϕ is logically true relative to a context $c \models_c \phi$ —if and only if ϕ is true in every model M relative to every coordinated variant of c in M , the occurrence tracking parameter 1, and assignment f :

$$\models_c \phi \text{ iff for every } M, c' \in M, \text{ and } f, \text{ if } R_{c'} = R_c, \text{ then } \models_{c'f1}^M \phi$$

It is easy to verify that \models_c (7) if and only if the two occurrences of ‘that person’ in (7) are coordinated relative to c .

Another example is due to Perry (1977), and developed by Braun (1996). Suppose that the bow and stern of the aircraft carrier USS Enterprise are visible to us, but the middle is hidden behind a tall building. Consider two uses of (11):

That is identical to that. (11)

¹⁴In this definition, I use

$$c \in M$$

as an abbreviation for

$$\exists C \exists D \exists I (M = \langle C, D, I \rangle \ \& \ c \in C)$$

On one use of (11), I point twice to the bow of the USS Enterprise. On the other, I point first to the bow, and then to the stern. Though both uses of (11) are true, they have different epistemic properties. In the first case, my audience can see that my use of (11) is true merely in virtue of their competence with ‘that’ and their recognition of my referential intentions. In the second case, my audience requires further information: that I am pointing at the same ship with my two pointing gestures. This epistemic difference arises, in part, from a difference in their logical properties, just as the difference between ‘Hesperus is identical to Hesperus’ and ‘Hesperus is identical to Phosphorus’ is, in part, a logical difference.¹⁵ A logic of demonstratives should explain the different logical properties of these two uses of (11).

FOLD offers a natural account of the logical behavior of sentences like (7) and (11). (7) is a logical truth relative to some contexts but not others, depending on whether the occurrences of ‘that person’ in (7) are coordinated in the context. But it is precisely those uses of (7) in which the speaker intends the corresponding uses of ‘that person’ to refer to the same thing in the same way in which the speaker’s use of (7) has the force of a logical truth. Similar remarks hold for (11), and in fact (11) and (7) are logically true relative all and only the same contexts.¹⁶ Thus the proposed account of logical truth in Sect. 3 correctly predicts when a use of a sentence like (7) has the force of a logical truth.

Logical consequence is more subtle, because it introduces sequences of formulas. Focusing for a moment on single-formula consequence (for simplicity), we must choose between the following analyses (using ‘ \rightarrow ’ for a conditional in our metalanguage, and omitting the restriction ‘ $c' \in M'$ ’):

- (α) $\phi \models_c \psi$ iff for every M, c' , and f , ($R_{c'} = R_c \rightarrow (\models_{c'f1}^M \phi \rightarrow \models_{c'f1}^M \psi)$)
 (β) $\phi \models_c \psi$ iff for every M, c' , and f , ($R_{c'} = R_c \rightarrow (\models_{c'f1}^M \phi \rightarrow \models_{c'fs(\phi,1)}^M \psi)$)

The two proposals differ over whether, in our evaluation of ψ , we take into account the occurrences of ‘that’ in ϕ . According to (α), we do not—the occurrence tracking parameter resets to 1 in our evaluation of ψ . According to (β), we do—in (β), we evaluate ψ relative to the occurrence tracking parameter $s(\phi, 1)$. Evidence in favor of (β) comes from our judgments about arguments like (9), therefore (10)⁷ from Sect. 2:

¹⁵I am *not* claiming that Frege’s puzzle or the phenomenon of cognitive significance is merely a matter of logic. The logical difference between the two uses of (11) is *evidence* for a difference in cognitive significance between the two uses. But to identify a logical difference between them is not to give an account of the cognitive significance of either.

¹⁶There are different ways to approach this feature of utterances theoretically. One approach is to take some kind of common-reference intention to be a primitive feature of certain linguistic actions. A different approach is to identify an intention with some descriptive content that fixes the reference of a use of a demonstrative, and to specify that two uses are coordinated if and only if the descriptive content of the reference fixing intentions are the same. For the purposes of logic, what matters is that coordination occurs somehow.

Nothing is taller than that. (9)

So, that is not taller than that. (10)

Suppose, again, that when I utter (9), I point at Mt. Whitney, and that when I utter (10), I point first at Lone Pine Peak, and then again at Mt. Whitney. In this example, I have argued well. You can follow my reasoning and recognize its validity. What you must follow in the example is my use of gestures, and perhaps intonation and focus, that indicate coordination relations between my uses of ‘that’. In the example in the introduction, the occurrence of ‘that’ in (9) is coordinated with the second occurrence of ‘that’ in (10).

This suggests that in evaluating arguments for logical consequence, we must take into account the occurrences of ‘that’ in the premises in our evaluation of the conclusion. Let (9*) and (10*) be suitable translations of (9) and (10) into FOLD:

$$\forall x \sim T x \delta \quad (9^*)$$

$$\sim T \delta \delta \quad (10^*)$$

And let c be the following context:

⟨(Mt. Whitney, Lone Pine Peak, Mt. Whitney), R ⟩

Finally, we suppose that 1 bears R to 3. Now according to (α), the argument $\lceil 9^* \rceil$ therefore $\lceil 10^* \rceil$ is not valid relative to c . It is true that nothing around Lone Pine, CA is taller than Mt. Whitney. It is false that Mt. Whitney is not taller than Lone Pine Peak. In resetting the occurrence-tracking parameter to 1 in our evaluation of (10*), we misinterpret the conclusion. In contrast, definition (β) predicts that the argument $\lceil 9^* \rceil$ therefore $\lceil 10^* \rceil$ is valid relative to c , because the correct occurrences of ‘that’ are coordinated, and the correct values are assigned to the occurrences of ‘that’ in (10*). Insofar as logic should yield correct predictions about good and bad reasoning, we should prefer definition (β) of logical consequence for FOLD.

The proposal in this section further clarifies my objection, in the previous section, to Kaplan’s theory of logical truth and consequence. I do not object to the claim that logical properties require generalizing over contexts. Rather, I object to the claim that this is the *only* role for contexts in logic. We may now articulate an alternative: contexts also determine the coordination relations held fixed when we generalize. This argument turns on an important fact about contexts: some features of contexts reflect ways in which discourse and reasoning are structured over time.¹⁷ One way in which this temporal structure unfolds is in the objects introduced by demonstratives into, and referred to by demonstratives during, a conversation. In

¹⁷Zardini offers a sophisticated logical treatment of one aspect of this temporal structure (Zardini, 2014). The dynamic structure of discourse may be semantically significant in more than one way.

semantics, it has long been recognized that demonstratives differ from indexicals in that distinct occurrences of the same demonstrative may be assigned distinct values. Recognition that this reflects a temporal dimension of context is more recent. As Brian Rabern and Brian Pickel have emphasized to me in conversation, a semantics for multiple occurrences of demonstratives has to be dynamic. I agree, to at least the minimal extent reflected in this section.¹⁸

To summarize the lessons of this and the previous section: (i) The contexts of Kaplan's logic, and of ours, model certain features of the use of language in reasoning and conversation. (ii) Modeling our reasoning with demonstratives requires more of contexts than simply a sequence of objects to serve as the values of distinct occurrences. We also recognize when, in reasoning and conversation, two occurrences of a demonstrative are coordinated, and our judgments about the success of our reasoning with demonstratives sometimes depend on this. (iii) Kaplan's theory, on its recent dynamic interpretations, offers us the resources to represent this feature of reasoning and conversation in the contexts of our logic.¹⁹ (iv) The result yields an intuitive but rigorous account of how some arguments in a first-order logic containing demonstratives may be logically valid relative to some contexts, but not valid relative to others.²⁰

4 The Coordination in Context Assumption

One objection to the view developed in this paper is that logical consequence is supposed to be at least both necessary and formal, but the consequence relation defined in (β) is not formal in the sense required for logic.²¹ What formality amounts to, however, is not always clear. Beall and Restall (2006, ch. 2.5) cite

¹⁸This is the same thing as allowing context-shifts, or 'mid-argument context change' (Zardini, 2014). Braun was the first to move in this direction in the formal semantics of demonstratives (Braun, 1996).

¹⁹I am now inclined to recognize cases of failed coordination—contexts c such that R_{cij} and $d_i \neq d_j$, and to take the requirement on contexts ruling out such cases as more akin to a restriction to proper contexts.

²⁰Gauker rejects this approach to the problem of referential promiscuity:

One could try relativizing validity to context, and hold that in some contexts ['this is bigger than that; therefore, this is bigger than that'] is valid and in others it is not. But this runs contrary to our expectation that while the truth value of a sentence may be relative to context, the logical properties of a sentence are fixed. In any case, it would be worthwhile to consider whether that presumption can be preserved. (Gauker, 2014, 292)

For the logic of demonstratives, I think that our examples, and Gauker's, show that this expectation is intuitively unfounded. I also find the cost of his alternative—each use of 'that' to refer to a new object is a new lexical item in an always growing language—to be excessive relative to the simplicity of the present approach. See Footnote 1.

²¹Sometimes being normative is added as a necessary condition for logical consequence as well.

MacFarlane's (2000) three different ways of understanding the claim that logic is formal, and end their discussion of formality "inconclusively" (Beall and Restall, 2006, 22). Yet as Glanzberg (2015, 77) notes, one traditional way of spelling out formality, associated with Tarski, is to identify a set of logical constants whose meaning is held fixed as models of the language vary. One lesson of the previous section is that we may (and should) understand demonstratives as formal in at least this regard. So any objection to the view of the present paper based on an appeal to formality must introduce some additional constraints on, or some alternative theory of, formality. In this section, I identify an assumption about logical form which, if true, is inconsistent with the theory of consequence developed in this paper. In the next section, I consider and reject one argument for this assumption.

To introduce this assumption, consider for a moment any quantifier-free formula of any first-order language, including FOLD. Let the formula be as propositionally complex as you would like. Now remove all of the constants and variables from the formula, leaving syntactic structure and predicates of various arities. One may identify, in this latter structure, the argument places vacated when we removed all constants and variables.

When two argument places in a sentence or argument are occupied by the same variable or non-demonstrative constant, then those two argument places are *coordinated*, where coordination amounts to a kind of semantic guarantee of coreference.²² The coordination of argument places imposed by the multiple occurrences of a variable or non-demonstrative constant is sometimes relevant to the validity of an argument, or the logical truth of a sentence or formula. The formula ' $a = a$ ' is a logical truth (assuming that '=' is a logical constant), while ' $a = b$ ' is not.

According to what I will call the *coordination assumption*, this is the only way for argument places to be coordinated in logic:

COORDINATION ASSUMPTION

The conclusion of an argument is a logical consequence of the premises of the argument, or a formula is a logical truth, only if the logical form of the argument or formula determines all logically relevant coordination relations between argument places in the argument or formula.

Yet according to the theory of consequence developed in this paper, some logically relevant coordination relations between argument places are determined in context. The syntax of a sentence or argument alone does not determine all coordination relations. Thus given the coordination assumption, the consequence relation defended in the previous section is not formal, in the sense required for logic, and whatever the investigation above, it is not an investigation into logical consequence.

Strictly, however, the proposal in the present paper is consistent with the coordination assumption as stated above. On the current proposal, (10) is a logical consequence of (9) only relative to some contexts but not relative to others. Thus it

²²This use of 'coordination' was introduced by Fine (2003, 2007) in his discussion of what he calls the 'antinomy of the variable'.

is strictly speaking correct that (10) is not a consequence of (9), when this claim is understood as entailing that (10) is always a consequence of (9). Thus on the present proposal, we may reject the antecedent of the coordination assumption. To generate an objection to the present proposal, we must modify the assumption.

To explain the required modifications, we may first distinguish the present view from an alternative recently defended by Iacona (2013, 2018). Iacona calls attention to examples, like the ones in this paper, of the use of demonstratives in reasoning, in which one argument can be valid on one occasion of use, but invalid on another. On his view, these arguments show that logical consequence, or validity, obtains or fails to obtain only relative to what he calls interpretations.

Thus far, Iacona's view may appear to be a variation on the proposal in this paper. Yet according to Iacona, an argument is valid relative to one interpretation, and invalid relative to another, only if that argument has different logical forms relative to the different interpretations. As a result, Iacona rejects what he calls the *intrinsicity presumption* about logical form: that the logical form of an argument as we use it in reasoning is intrinsic to the syntax of the argument we use.²³ Thus on Iacona's theory, it is not just the validity of an argument, but also its logical form, that varies from one interpretation to another.

According to Iacona, we represent the logical form of a sentence or argument by translating that sentence or argument into a privileged formal language designed for the study of logical form. To illustrate, consider again our original example of an argument containing multiple occurrences of a demonstrative:

Nothing is taller than that. (9)

So, that is not taller than that. (10)

According to Iacona, in a situation in which we use this argument to reason validly, the logical form of this argument may be represented by the first order formulas (12) and (13):

$\forall x \sim Txa,$ (12)

$\sim Tba.$ (13)

But we might also reason invalidly, if (for example) the logical form of the argument as we use it is represented by (12) and (14):

²³Beall and Restall (2006, 20–21) appear to endorse something like the intrinsicity presumption:

The *form* of a thought is perhaps best construed as the structural features intrinsic to the propositional content, rather than any of its accidental features.

$$\forall x \sim Txa, \tag{12}$$

$$\sim Tbc. \tag{14}$$

In this way, our argument has a different logical form in the different contexts in which we use it. Since the syntax of the argument we use remains unchanged across the examples, Iacona rejects the intrinsicality presumption.²⁴

On Iacona's theory, the logical form of an argument, relative to a context c , determines all logically relevant coordination relations because the logical form of the argument, relative to the context, contains no referentially promiscuous expressions. All logically relevant coordination relations are determined by the syntax of the formulas used to illustrate the logical form of a sentence or argument. Thus on Iacona's view, the formulas (9*) and (10*) of FOLD are not allowed in the privileged formal language used to represent logical form:

$$\forall x \sim Tx\delta \tag{9*}$$

$$\sim T\delta\delta \tag{10*}$$

Instead, on Iacona's view arguments in English containing referentially promiscuous expressions will be assigned distinct logical forms relative to distinct contexts.

For Iacona, an argument is valid relative to a context c only if its logical form in c syntactically determines all logically relevant coordination relations. This feature of Iacona's view shows how to avoid the earlier problem with the coordination assumption. We thus arrive at our modified assumption:

COORDINATION IN CONTEXT ASSUMPTION

The conclusion of an argument is a logical consequence, relative to a context c , of the premises of the argument only if the logical form of the argument, relative to c , syntactically determines all logically relevant coordination relations between argument places in the argument.

The antecedent of the coordination in context assumption allows that an argument may be logically valid relative to some contexts but not others. Thus the coordination in context assumption avoids the immediate response to the coordination assumption above.

The coordination in context assumption is inconsistent with the theory of logical consequence developed in the present paper. On the theory of the present paper, an argument may be valid relative to a context c , even though the logical form of that argument, because it contains demonstratives, does not syntactically encode all

²⁴It is not essential to the rejection of the intrinsicality presumption that one follow Iacona in using constants. We might instead use free variables, since distinct occurrences of the same free variable also impose a coordination relation on the argument places at which the variable occurs. The use of variables was suggested to me by Murali Ramachandran, whom I thank for spirited discussion of the view in this paper.

logically relevant coordination relations. Thus if formality for logical consequence requires that all logically relevant coordination relations are syntactically determined at the level of logical form, then the present theory of logical consequence fails as a formal theory.

The formal tools of the present paper, however, suggest an alternative to the coordination in context assumption. To bring out this alternative, we may follow Iacona in taking the logical form of an argument, relative to a context, to be revealed by translating that argument into a privileged formal language. But unlike Iacona, we will take our formal language to contain referentially promiscuous expressions. Indeed, we may use FOLD itself. We may now say that the logical form of an argument, relative to a context c , determines all logically relevant coordination relations because the context c determines at least some logically relevant coordination relations. In this way, we may retain what Iacona calls the intrinsicality presumption, and assign to an argument like (9), therefore (10), in any context, the logical form given by the formulas (9*) and (10*). This logical form does not change from one context to another. What changes from one context to another are the coordination relations between the occurrence of ‘ δ ’ in the argument.

This shows that we may not simply assume the coordination in context assumption in objecting to the theory of the present paper. We now know that we can incorporate demonstratives directly into first-order languages, without any changes to the syntax of those languages. Yet what distinguishes demonstrative-free first order languages is that all logically relevant coordination relations between argument places are fully determined by syntax. So to restrict the formal languages one uses to represent logical forms to demonstrative-free first-order languages, as Iacona does, is just to beg the question in favor of the coordination in context assumption.²⁵

The argument of Sect. 3 shows that the logic of demonstratives may be thought of as formal in at least the standard model-theoretic sense introduced by Tarski, because we may treat demonstratives as logical constants. As a result, it is consistent with this Tarskian conception of formality that we reject the coordination in context assumption. Thus to generate an objection to the theory of consequence developed in this paper, we must provide an independent argument for this assumption. In the next section, I consider, and reject, one such argument.

5 Demonstratives and Proof

According to the coordination in context assumption, the formulas we use to represent the logical form of natural language sentences and arguments, relative to contexts, do not contain any referentially promiscuous expressions. Thus on the

²⁵See my review of Iacona’s *Logical Form: Between Logic and Natural Language*, forthcoming in *Dialectica*, for further discussion of this point (Georgi, 2019).

coordination in context assumption, if a natural language argument is valid, relative to a context c , then the formula that represents the logical form, relative to c , of the conclusion of the argument, is a logical consequence, full stop, of the formulas that represent the logical forms, relative to c , of the premises of the argument. At the level of logical form, according to the coordination in context assumption, logical consequence behaves as it has always behaved. That some arguments of natural language are valid relative to some contexts but not others does not reveal anything deep about logic. Rather, it shows why referential promiscuity is something to be avoided in the formal languages we use to represent logical forms.

One argument for this picture is based on the soundness and completeness of standard systems of proof for first-order languages. A system of proof S is sound if and only if, if ϕ is provable in S from Γ then ϕ is a consequence of Γ . A system of proof S is complete if and only if, if ϕ is a consequence of Γ then ϕ is provable in S from Γ .²⁶ Now let S be a sound and complete system of proof for the privileged formal language we use to represent logical forms. The argument proceeds as follows: first, we observe that a formula or argument some of whose logically relevant coordination relations are undetermined may, in some situations, be false or invalid. Thus if S is sound, then any proof in S must terminate in a formula or argument all of whose logically relevant coordination relations are determined. Yet since provability is a fundamentally syntactic notion, a proof terminates in a formula or argument all of whose logically relevant coordination relations are determined only if those coordination relations are determined by the syntax of the formula or argument. Finally, if S is complete, it then follows that the consequence relation obtains only between formulas all of whose logically relevant coordination relations are determined by syntax. Since the consequence relation in question obtains between the formulas of the privileged formal language we use to represent logical forms, we now have an argument for the coordination in context assumption. In what follows, I will call this the *argument from provability*.

I have two objections to the argument from provability. To explain these objections, I will first introduce two different ways of incorporating contexts into proofs. For concreteness, I will focus on systems of sequent rules. Following Gentzen (1964, 290), we understand a sequent $\Delta \Rightarrow \Gamma$, where Δ and Γ are sequences of formulas, as equivalent to a material conditional whose antecedent is the conjunction of the members of Δ (in order), and whose consequent is the disjunction of the members of Γ (also in order).²⁷ As a simple example, let us take a simple first-order sequent calculus—call it DL—including a rule for thinning the premises of a sequent:

²⁶Blanchette offers a nuanced discussion of the philosophical significance of such basic formal results (Blanchette, 2001).

²⁷Some recent discussions of the sequent calculus use multisets of formulas in place of sequences of formulas (Beall et al., 2018). Multisets are like sequences in that they may contain multiple occurrences of a single element, but they are like sets in that they do not impose any order on their members. For demonstratives, however, this approach will not do. Changing the order of formulas in a sequence will change the coordination relations between any demonstratives occurring in those formulas.

$$\frac{\phi_1, \dots, \phi_n \Rightarrow \psi_1, \dots, \psi_m}{\phi_1, \dots, \phi_n, \phi_{n+1} \Rightarrow \psi_1, \dots, \psi_m}$$

One way to incorporate contexts into the sequent calculus is to relativize the sequents themselves to contexts:

$$\frac{\phi_1, \dots, \phi_n \Rightarrow_c \psi_1, \dots, \psi_m}{\phi_1, \dots, \phi_n, \phi_{n+1} \Rightarrow_{c'} \psi_1, \dots, \psi_m}$$

Given our understanding of sequents, the semantics for this proposal is straightforward. A sequent relative to a context c is equivalent to the corresponding conditional relative to c . Call the resulting system of sequent rules DL1.

One result of relativizing sequents to contexts is that for each rule of DL1, we must also specify how the context of the output sequent is related to the contexts of the input sequents. The shift from c to c' in the thinning rule above, for example, reflects the obvious fact that if ϕ_{n+1} contains occurrences of ‘ δ ’, then adding it to the output sequent will affect the coordination relations between occurrences of ‘ δ ’ in the formulas of the input sequent.²⁸ Other rules, such as the rule for introducing a conditional to the right of a sequent, will not require any change in the context:

²⁸ To fully flesh out the thinning rule, we must specify how c' is related to c . The details are a little messy, but it is easy to see in outline how to proceed: Let k be the number of occurrences of ‘ δ ’ in the sequence of premises of the input sequent to an application of thinning, let l be the number of occurrences of ‘ δ ’ in ϕ_{n+1} —the formula being added by conjunction. We first characterize the coordination relation $R_{c'}$, in three stages. I will introduce and reflect on the first two before I turn to the third. First, for any positive integers i and j such that $i, j \leq k$, $R_{c'}ij$ if and only if R_cij . Second, for any positive integers i and j such that $i \leq k$ and $j > (k+l)$, $R_{c'}i(j-l)$ if and only if R_cij . At this point, $R_{c'}$ simply ignores the occurrences of ‘ δ ’ in ϕ_{n+1} .

But we should also permit occurrences of ‘ δ ’ in ϕ_{n+1} to be coordinated with occurrences of ‘ δ ’ in either ϕ_1 through ϕ_n or ψ . In the third stage, where we introduce new coordination relations, we assume that an occurrence of ‘ δ ’ in ϕ_{n+1} is coordinated with any two occurrences of ‘ δ ’ in ϕ_1 through ϕ_n or ψ , only if those latter two occurrences are already coordinated at the second stage. The following two applications of thinning illustrate this (restricting attention to only the coordination relations):

$$\text{thinning 1: } \frac{F\delta \Rightarrow_{\{(1,2)\}} F\delta}{F\delta, G\delta \Rightarrow_{\{(1,3)\}} F\delta}$$

$$\text{thinning 2: } \frac{F\delta \Rightarrow_{\{(1,2)\}} F\delta}{F\delta, G\delta \Rightarrow_{\{(1,2), (1,3), (2,3)\}} F\delta}$$

In thinning 1, we choose not to let the occurrence of ‘ δ ’ in ‘ $G\delta$ ’ be coordinated with any other occurrences of ‘ δ ’ in the output sequent. In thinning 2, we choose to let the occurrence of ‘ δ ’ in ‘ $G\delta$ ’ be coordinated with the other, already coordinated, occurrences of ‘ δ ’.

$$\frac{\phi_1, \dots, \phi_n, \phi \Rightarrow_c \psi, \psi_1, \dots, \psi_m}{\phi_1, \dots, \phi_n \Rightarrow_c (\phi \supset \psi), \psi_1, \dots, \psi_m}$$

In this rule, the sequences of occurrences of ‘ δ ’ is unchanged from the input sequent to the output sequent. As a result, no changes are required to the context by an application of this rule.

In practice, in constructing proofs in DL1, what matters is the construction of the coordination relation. Thus to illustrate the system DL1, we may construct a proof of

$$(F\delta \supset (G\delta \supset F\delta)),$$

relative to a context c whose coordination relation includes $\{\langle 1, 3 \rangle\}$. The proof is straightforward:

$$\frac{\frac{F\delta \Rightarrow_{\{\langle 1, 2 \rangle\}} F\delta}{F\delta, G\delta \Rightarrow_{\{\langle 1, 3 \rangle\}} F\delta}}{F\delta \Rightarrow_{\{\langle 1, 3 \rangle\}} (G\delta \supset F\delta)}}{\Rightarrow_{\{\langle 1, 3 \rangle\}} (F\delta \supset (G\delta \supset F\delta))}$$

This proof starts with an instance of what is sometimes called the identity axiom:

$$\frac{}{\phi \Rightarrow_c \phi}$$

In DL1, we stipulate that if ϕ contains j occurrences of ‘ δ ’, then R_c contains as a subset at least $\{\langle 1, j + 1 \rangle, \langle 2, j + 2 \rangle, \dots, \langle j, j + j \rangle\}$. In the application of thinning, we ignore the occurrence of ‘ δ ’ in ‘ $G\delta$ ’ (see Footnote 28). Two applications of our rule for introducing a conditional to the right of a sequent do not affect the coordination relation.

Another way we might try to incorporate contexts into the sequent calculus is to relativize each formula in the sequent to a context (I write ‘ (ϕ, c) ’ for ‘ ϕ relative to c ’):

$$\frac{(\phi_1, c_1), \dots, (\phi_n, c_n) \Rightarrow (\psi_1, c_{n+1}), \dots, (\psi_m, c_{n+m})}{(\phi_1, c_1), \dots, (\phi_n, c_n), (\phi_{n+1}, c) \Rightarrow (\psi_1, c_{n+1}), \dots, (\psi_m, c_{n+m})}$$

Call the resulting system of sequent rules DL2. Unlike DL1, the sequents of DL2 cannot be interpreted using a material conditional of the object language (without defining some further context out of the individual contexts introduced in the sequent). We may instead use a material conditional in the metalanguage: if ϕ_1 is true relative to c_1 , and \dots , and ϕ_n is true relative to c_n , then either ψ_1 is true relative to c_{n+1} , or \dots , or ψ_m is true relative to c_{n+m} .

The system DL1 is a natural proof system for the analysis of logical consequence developed in this paper.²⁹ Thus I propose to set DL2 aside.³⁰ DL1 is interesting enough. For example, it admits of basic soundness and completeness results:

Theorem 1 (Soundness and Completeness for DL1) *A sequent $\Delta \Rightarrow_c \Gamma$ is provable in DL1 if and only if Γ is a consequence of Δ relative to c .*³¹

I here only sketch the proofs. For soundness, we have to specify a feature of the basic sequent rule $\phi \Rightarrow_c \phi$. For any application of the rule, the formula ϕ must contain some finite number n_ϕ of occurrences of ‘ δ ’. The restriction is then as follows: the complete rule is that $\phi \Rightarrow_c \phi$ only if for any $1 \leq i \leq n_\phi$, $R_c i(i + n_\phi)$. It is trivial

²⁹For a final example, consider a rule for introducing a universal quantifier into the premises of a sequent:

$$\frac{\phi_1, \dots, \phi_n, \phi(t) \Rightarrow_c \psi_1 \dots \psi_m}{\phi_1, \dots, \phi_n, \forall x \phi(x) \Rightarrow_{c'} \psi_1 \dots \psi_m}$$

For this rule, we require that all occurrences of ‘ δ ’ replaced by variables in an application of the rule be coordinated. The shift from c to c' in this rule tracks the occurrences of ‘ δ ’ replaced by variables. Again, the details are messy, but the underlying goal is clear—intuitively, we remove from the coordination scheme any pairs of integers corresponding to occurrences of ‘ δ ’ that have been replaced, and then shift the remaining elements of the coordination scheme to reflect the reduction in the total number of occurrences of ‘ δ ’. The following proof illustrates the rule:

$$\frac{\sim T \delta \delta \Rightarrow_{\{(1,3), (2,4)\}} \sim T \delta \delta}{\forall x \sim T x \delta \Rightarrow_{\{(1,3)\}} \sim T \delta \delta}$$

This is a proof of the argument considered in the introduction relative the coordination scheme of the context introduced toward the end of Sect. 3. The coordination scheme $\{(1, 3)\}$ is the result of first removing the pair $\langle 1, 3 \rangle$ from the coordination scheme of the input sequent, and then subtracting 1 from all the integers in the remaining coordination scheme.

³⁰Radulescu (2015) offers a natural logical interpretation for the sequents of DL2. A sequent $(\phi_1, c_1), \dots, (\phi_n, c_n) \Rightarrow (\psi, c)$ or argument

$$\boxed{\phi_1 \text{ relative to } c_1, \dots, \phi_n \text{ relative to } c_n; \text{ therefore } \psi \text{ relative to } c_{n+1}} \quad (15)$$

determines a sequence of contexts

$$\langle c_1, \dots, c_{n+1} \rangle. \quad (16)$$

Radulescu defines an equivalence relation of similarity on sequences of contexts, and then proposes that (15) is valid if and only for every sequence of contexts

$$\langle c'_1, \dots, c'_{n+1} \rangle$$

similar to (16), if ϕ_1 is true relative to c'_1, \dots , and ϕ_n is true relative to c'_n , then ψ is true relative to c'_{n+1} .

³¹The definition of consequence in Sect. 3 is stated only for single consequences, but I see no reason not to extend it to multiple consequences, where these are interpreted in a specific order. Thus for the purposes of discussing the proof theory, I assume this generalization.

in this case that if $\phi \Rightarrow_c \phi$ is provable then $\phi \models_c \phi$. We then prove that each of the sequent rules preserves this property.

For completeness, we proceed less directly. (I have yet to find a more direct proof.) Suppose that Γ is a consequence of Δ relative to c . We construct a sequent $\Delta' \Rightarrow \Gamma'$ that contains no occurrences of ' δ ', and hence which if provable is provable using rules that do not require any relativization to context. The construction of the demonstrative-free sequent proceeds by replacing all coordinated occurrences of ' δ ' with the same constant (and using a distinct constant for each equivalence class of coordinated occurrences). We then argue by induction that Γ is a consequence of Δ relative to c if and only if Γ' is a consequence of Δ' . Since the latter consequence does not turn on the semantics of demonstratives, traditional completeness results entail that the sequent $\Delta' \Rightarrow \Gamma'$ is provable. Given a proof of this sequent, we can construct a proof of $\Delta \Rightarrow_c \Gamma$ by substitution of demonstratives for the appropriate constants and the construction of the appropriate coordination schemes.

Theorem 1 reveals one flaw with the argument from provability. According to the argument, since provability is a fundamentally syntactic notion, a proof terminates in a formula or argument all of whose logically relevant coordination relations are determined only if those coordination relations are determined by the syntax of the formula or argument. We now see that it is possible to devise a system of proof such that at least some logically relevant coordination relations, in provable arguments and formulas, are not determined syntactically. DL1 is just such a system.

The techniques used to prove Theorem 1 can also be used to prove a related result for the unrelativized sequent rules of DL, applied to the language of FOLD:

Theorem 2 (Soundness and Completeness for DL) *A sequent $\Delta \Rightarrow \Gamma$ (where Δ and Γ may contain occurrences of ' δ ') is provable in DL if and only if there is some context c such that Γ is a consequence of Δ relative to c .*

Again, I sketch the proofs. For soundness, suppose that a particular sequent $\Delta \Rightarrow \Gamma$ is provable. Substitute for each step of this proof the corresponding context-relativized sequent rule. The result will be a proof, for some c , of the sequent $\Delta \Rightarrow_c \Gamma$ in the context-relativized sequent calculus. By Theorem 1 Γ is a consequence of Δ relative to c .

For completeness, assume that there is some context c relative to which Γ is a consequence of Δ . We generate the ' δ '-free sequences Δ' and Γ' by substituting constants for occurrences of ' δ ' in accordance with R_c , and prove that Γ' is a consequence of Δ' . Then by the completeness of standard first-order logic, the sequent $\Delta' \Rightarrow \Gamma'$ is provable. Now a substitution of ' δ ' for the appropriate constants yields a proof of the sequent $\Delta \Rightarrow \Gamma$.

Theorem 2 reveals a second flaw with the argument from provability. According to the argument, if a system of proof S is sound, then any proof in S must terminate in a formula or argument all of whose logically relevant coordination relations are determined. We already know, from Theorem 1, that these coordination relations do not need to be determined by the syntax of the formula or argument. Yet Theorem 2 reveals that a proof in DL can terminate in a formula or argument of FOLD some of

whose logically relevant coordination relations are not determined at all. The key to this result is to understand what it shows about the formula or argument at which a proof terminates. According to Theorem 2, the existence of such a proof shows that there is some context relative to which the proven argument is valid, or the proven formula is a logical truth. The proof does not provide any immediate way to construct such a context (though for any proof in DL there will be at least one corresponding proof in DL1 that does construct the required context).

The proofs of Theorems 1 and 2 invoke an unrelativized logical consequence relation. Does this not show that I am in fact committed to such a relation, despite my claim that logical consequence does or does not obtain only relative to a context? This concern is unfounded. For a first-order language containing no referentially promiscuous expressions, a traditional treatment of semantics and logic (I include Kaplan in this tradition) is sufficient for studying the logical properties of the language. The logical properties of such demonstrative-free languages have been studied extensively. But for any language that does contain a referentially promiscuous expression, logical consequence obtains or fails to obtain only relative to contexts. Assuming that the English demonstrative pronoun ‘that’ is a referentially promiscuous expression, it follows that logical consequence in English, if there is any such thing, obtains or fails to obtain only relative to contexts.

6 Conclusion

The first-order logic of demonstratives shows us that contexts can play surprising roles in logic. I have argued in this paper that for first-order languages containing demonstratives or other referentially promiscuous expressions, logical consequence and logical truth obtain or fail to obtain only relative to contexts. I have shown how to implement this idea using the formal tools developed by Kaplan and Braun. The result is, I suggest, a simple and intuitive account of the role of demonstratives in first-order logic.

The resulting theory of logical consequence is inconsistent with at least one interpretation of a principle we called the coordination in context assumption. According to this interpretation, demonstratives, and referential promiscuity generally, are seen as something to be eliminated from our logical investigations as much as possible, and especially from the formal languages we use to represent logical forms.

As an objection to the theory of consequence in the present paper, however, inconsistency with the coordination in context assumption is only as strong as the strength of our reasons to endorse the assumption. In this paper, I have shown that basic metalogical results concerning the first-order logic of demonstratives undermine one potentially significant argument in favor of the coordination in context assumption. If anything, the metalogical results in Sect. 5 should move us to reject the excessively syntactic coordination in context assumption, and to recognize that some logically relevant coordination relations are determined by the contexts in which we reason with demonstratives.

Acknowledgements I want to thank Michael Glanzberg, Murali Ramachandran, and the participants of Logica 2017 and of the First Context, Cognition and Communication conference for suggestions and discussion that led to improvements of this paper. I also want to thank Tadeusz Ciecierski and Paweł Grabarczyk for their work as editors of the present volume.

References

- Beall, J. C., & Restall, G. (2006). *Logical pluralism*. Oxford: Clarendon Press.
- Beall, J. C., Glanzberg, M., & Ripley, D. (2018). *Formal theories of truth*. New York: Oxford University Press.
- Blanchette, P. A. (2001). Logical consequence. In L. Goble (Ed.), *The Blackwell guide to philosophical logic* (pp. 115–135). Malden: Blackwell Publishing.
- Braun, D. (1996). Demonstratives and their linguistic meanings. *Noûs*, 30(2), 145–173.
- Braun, D. (2008). Complex demonstratives and their singular contents. *Linguistics and Philosophy*, 31, 57–99.
- Elbourne, P. (2008). Demonstratives as individual concepts. *Linguistics and Philosophy*, 31, 409–466.
- Fine, K. (2003). The role of variables. *Journal of Philosophy*, 80(12), 605–631.
- Fine, K. (2007). *Semantic relationism*. Malden: Blackwell Publishing.
- Gauker, C. (2014). How many bare demonstratives are there in English? *Linguistics and Philosophy*, 37, 291–314.
- Gentzen, G. (1964). Investigations into logical deduction. *American Philosophical Quarterly*, 1(4), 288–306.
- Georgi, G. (2012). Reference and ambiguity in complex demonstratives. In W. Kabasenche, M. O'Rourke, & M. H. Slater (Eds.), *Reference and referring: Topics in contemporary philosophy* (Vol. 10, pp. 357–384). Cambridge, MA: MIT Press.
- Georgi, G. (2015). Logic for languages containing referentially promiscuous expressions. *Journal of Philosophical Logic*, 44(4), 429–451.
- Georgi, G. (2019). Review of *Logical form: Between logic and natural language*, by Andrea Iacona. *Dialectica*, 72(4), 622–628.
- Glanzberg, M. (2015). Logical consequence and natural language. In C. R. Caret & O. T. Hjortland (Eds.), *Foundations of logical consequence*. Corby: Oxford University Press.
- Hawthorne, J., & Manley, D. (2012). *The reference book*. Oxford: Oxford University Press.
- Iacona, A. (2013). Logical form and truth conditions. *Theoria*, 78, 439–457.
- Iacona, A. (2018). *Logical form: Between logic and natural language*. Synthese Library. Cham: Springer International Publishing.
- Kamp, H. (1971). Formal properties of 'now'. *Theoria*, 37, 237–273.
- Kaplan, D. (1979). On the logic of demonstratives. *Journal of Philosophical Logic*, 8(1), 81–98.
- Kaplan, D. (1989). Demonstratives. In J. Almog, J. Perry, H. K. Wettstein, & D. Kaplan (Eds.), *Themes from Kaplan*. New York: Oxford University Press.
- King, J. C. (2001). *Complex demonstratives: A quantificational account*. Cambridge, MA: MIT Press.
- Lewis, D. (1970). General semantics. *Synthese*, 22(1/2), 18–67.
- Lewis, D. (1981). Index, context, and content. In S. Kanger & S. Ohman (Eds.), *Philosophy and grammar*. Dordrecht: D. Reidel.
- MacFarlane, J. (2000). *What does it mean to say that logic is formal?* Ph.D. thesis, University of Pittsburgh.
- Montague, R. (1974). *Formal philosophy: Selected papers of Richard Montague*. New Haven: Yale University Press.
- Perry, J. (1977). Frege on demonstratives. *Philosophical Review*, 86(4), 474–497.

- Pickel, B., Rabern, B., & Dever, J. (2018). Reviving the parameter revolution in semantics. In D. Ball & B. Rabern (Eds.), *The science of meaning* (pp. 138–171). Oxford: Oxford University Press.
- Predelli, S. (2005). *Contexts: Meaning, truth, and the use of language*. Oxford: Clarendon Press.
- Radulescu, A. (2015). The logic of indexicals. *Synthese*, 192(6), 1839–1860.
- Salmon, N. (2002). Demonstrating and necessity. *Philosophical Review*, 111(4), 497–537.
- Seegerberg, K. (1973). Two-dimensional modal logic. *Journal of Philosophical Logic*, 2(1), 77–96.
- Stalnaker, R. (1972). Pragmatics. In G. Harman & D. Davidson (Eds.), *The semantics of natural languages*. Dordrecht: Reidel.
- Stalnaker, R. (2014). *Context*. Oxford: Oxford University Press.
- Sullivan, A. (2013). *Reference and structure in the philosophy of language: A defense of the Russellian orthodoxy*. New York: Routledge.
- Van Orman Quine, W. (1951). *Mathematical logic* (Revised ed.). Cambridge, MA: Harvard University Press.
- Yagisawa, T. (1993). Logic purified. *Noûs*, 27(4), 470–486.
- Zardini, E. (2014). Context and consequence. An intercontextual substructural logic. *Synthese*, 191, 3473–3500.

De Se as Variable Binding: On Context-Sensitivity in Utterance Reports



Joan Gimeno-Simó

Abstract In this paper I will be defending a view of *de se* attitudes according to which all the inaccuracies attributed to Kaplan's Logic of Demonstratives having to do with the nature of *de se* attitude reports follow a single pattern: they are argued to correspond to bound uses of the variables contributed by the pronouns. I will argue that a traditional view of propositions suffices for dealing with these phenomena from a linguistic point of view.

Keywords Indexicality · Monsters · Self-ascription · Oratio obliqua · Logophoricity

1 Introduction

David Kaplan's highly influential theory of indexicals (Kaplan 1989a, b) is committed to a certain view on the interactions between intensionality and indexicality: namely, that indexicals are directly referential expressions, and as such they are not affected by the behavior of intensional (modal or temporal) operators: indexicals depend on context, but once their value is determined, it is fixed once and for all, since no operators can shift the context of utterance. A consequence of Kaplan's views on indexicality and content is that the so-called essential indexicality (Perry 1979) or *de se* nature (Lewis 1979) of self-ascriptive attitudes is systematically lost when such attitudes are reported: in Kaplan's terminology, what matters in attitude reports is the original utterance's content, and not its character, which amounts to saying that indirect discourse is sensitive only to the information transmitted by the reported utterance, regardless of the context in which it took place.

It is usual to present two distinct sets of data against Kaplan's ideas: on the one hand, the existence of indexical expressions that happen to be sensitive to

J. Gimeno-Simó (✉)

Department of Philosophy, University of València, València, Spain
e-mail: joan.gimeno@uv.es

the original utterance's character, such as logophoric pronouns (Castañeda 1966, 1967; Schlenker 2002, 2003) and English constructions involving PRO (Chierchia 1989), and on the other hand shiftable indexicals (Schlenker 2002, 2003; Anand and Nevins 2004). The aim of this paper is to argue that it is possible to identify a pattern common to all the cases in which indirect speech must be sensitive to the character of the original utterance: *de se* and *de re* readings of a sentence stem from the behavior of the variables found at the level of logical form (LF). Namely, I will be arguing that *de se* readings occur whenever the variables contributed by pronouns are bound. A consequence of this view is that the traditional resources of semantics are enough for handling the phenomenon of the *de se* in attitude reports.

This paper is structured in three parts. First, I will summarize Kaplan's views on indexicality and intensionality, as well as the consequences of his commitments to direct reference (Sect. 2). Next, I will discuss the above-mentioned data and show that it is possible to associate it with a certain sort of linguistic construction (Sects. 3 and 4). Finally (Sect. 5), I will show that, if my account is correct, it is possible to deal with the phenomenon of *de se* attitudes without resorting to properties nor to any non-conventional notion of proposition – even though a certain kind of monster operators is required.

2 Kaplan on Direct Reference

In his Logic of Demonstratives (henceforth LD), Kaplan (1989a, b) distinguished two kinds of meanings for expressions: their *content* and their *character*. The former is the customary truth-conditional contribution of an expression to the proposition expressed, whereas the second is a rule for determining its content in each context. In Kaplan's LD, as in other double-indexed systems, extensions of expressions are assigned with respect to two different parameters: a context c and a circumstance of evaluation i . A context is a tuple $\langle c_A, c_L, c_W, c_T \rangle$, where each parameter denotes, respectively, the agent, the location, the world and the time of c . Circumstances, on the other hand, are understood as tuples $\langle w, t \rangle$ consisting of a possible world w and a time t . An important thing to bear in mind is that, in assessing the extension of a sentence at a context and a circumstance, the time and possible world of the latter needn't coincide with the world and the time determined by the former (i.e., with c_W and c_T).¹ As usual in model-theoretic semantics, there is a further relativization to an assignment, so that sentences are true or false with respect to a context c , an index i and an assignment g . Assignments are usually dropped for pedagogical reasons when discussing Kaplan's LD, but, since they are important for the purposes

¹However, equating the parameters of the circumstance to those of the context yields a less relativized notion of truth, which is the one that Kaplan employs for issues of analyticity and logical validity (see Kaplan 1989a: 547).

of this essay, I won't omit them. Thus the extension of p with respect to a context c , an assignment g and a circumstance $\langle w, t \rangle$ may be represented as $[[p]]^{c, g, w, t}$.

In LD, indexicals are distinguished for having a variable character: whereas most expressions (proper names, definite descriptions, modal operators, verbs, adjectives, etc.) contribute the same content in each context, the character of indexicals picks out different individuals in different contexts. Thus, for instance, the character of "I" picks out, for each context c , c_A . More accurately, and given that Kaplanesque contents are understood as functions from circumstances to extensions, the character of "I" is a function that, for each context c , returns the function (content) that, for each circumstance i , determines c_A as the extension of "I" in i . The same holds *mutatis mutandis* for the character of "here", "now" and "actually".

All this entails that indexicals are *directly referential*: though their characters are variable functions, their contents are given by constant functions. This means that their semantic value is fixed once and for all by the context, and circumstances are irrelevant for assessing such value. Being directly referential amounts to excluding any Fregean *Sinn*, any mode of presentation, from the truth-conditional import (*content*) of an expression, so that the only thing that enters into content is simply the individual designated by that expression. In particular, it is worth noting that modal and temporal operators bear no effect on indexicals, i.e., their semantic value is unaffected by the possible world or time in which the sentence is being assessed:

- (1) $[[\text{Possibly, I am tired}]]^{c, g, w, t} = 1$ iff there is some world v such that $[[c_A \text{ is tired}]]^{c, g, v, t} = 1$
- (1a) $[[\text{Possibly the speaker is tired}]]^{c, g, w, t} = 1$ there is some world v such that $[[\text{The speaker is tired}]]^{c, g, v, t} = 1$
- (1b) $[[\text{Someday everyone now alive will be dead}]]^{c, g, w, t} = 1$ iff there is some day d such that for all a ($[[x \text{ is alive}]]^{c, g(a/x), w, c^T} = 1 \rightarrow [[x \text{ is dead}]]^{c, g(a/x), w, d} = 1$), where $g(a/x)$ is just like g except that it assigns a to x .
- (1c) $[[\text{It is possible for everything that is actually red to be shiny}]]^{c, g, w, t} = 1$ iff there is some world v such that for all a ($[[x \text{ is red}]]^{c, g(a/x), c^W, t} = 1 \rightarrow [[x \text{ is shiny}]]^{c, g(a/x), v, t} = 1$), where $g(a/x)$ is just like g except that it assigns a to x .

In (1), the indexical "I" picks out the speaker and rigidly designates her, which means that, in order to assess the semantic value of the sentence, it is not necessary to look at who is speaking in some other possible world: we just have to check whether there is some possible world in which a certain object, c_A , is tired. This contrasts sharply with (1a), in which we do have to look for speakers across possible worlds. In (1b) and (1c) the indexicals "now" and "actually" are respectively under the scope of the intensional operators "someday" and "it is possible", which shift the time and

world at which the whole sentence is being assessed, yet this does not prevent them from picking out their referents from the context of utterance.²

This brings us to another important point of Kaplan's commitment to direct reference: his ban on *monsters*. Operators like "possibly" or "someday" are respectively modal and temporal operators, i.e., operators that shift the world and time at which the sentence is being assessed. One may suspect that it is possible to introduce similar operators for shifting the context parameter. But such kind of operators, which Kaplan dubs "monsters", would contradict his ideas. Let us devise a monster operator: let M be an operator such that $[[Mp]]^{c, g, w, t} = 1$ iff, in some contexts k , $[[p]]^{k, g, w, t}$. Compare then (2) to (2a):

(2) I am tired

(2a) M (I am tired)

If (2) is uttered in context c , then its content is a function that assigns truth to every circumstance in which c_A is tired. Namely, its content may be represented as

$\lambda w. \lambda t. c_A$ is tired in w, t

But things are quite different with (2a), since an utterance of this sentence in context c would yield a content that has nothing to do with c_A . This is more easily seen if we forget for a moment about intensionality and think of propositions as structured entities (as Kaplan himself often does). (2b) is the structured proposition representing the content of (2) in c :

(2b) $\langle c_A, \text{being tired} \rangle$

On the other hand, a structured proposition representing the information conveyed by (2a) would look like this (Rabern and Ball 2019: 402–403):

(2c) $\langle \langle \text{the agent, being tired} \rangle, \text{being true at some contexts} \rangle$

Of course, (2c) has nothing to do with c_A . Allowing this kind of operators would entail abandoning Kaplan's views on direct reference. Namely, the reason why indexicals are directly referential is that their contribution to content is just an object. Now, content is an important notion in Kaplan, since it amounts to "what is said", "the information conveyed by an utterance" and "the object of assertion and thought": contents are the assertoric value of sentences in a context. What this means is that the information conveyed by a sentence like "Angela Merkel is tired" is the same as the content of an utterance of (2) by Angela Merkel, and, indeed, if uttered by Angela Merkel, both would be reported in the same way ("Angela Merkel

²Saying that these words "pick out a referent" is not completely accurate, since they are actually sentential operators that fix the parameters of the circumstance to the actual world and the current time. They are also treated as such in the papers from which these examples and discussion originated: Prior (1968) and Crossley and Humberstone (1977). Nothing relevant hinges on the use of this terminology.

said that she is tired”). This is so because a verb like “say” operates on contents, not on characters. Indeed, a report of the form “Angela Merkel said that I am tired” is not appropriate for this case.

This led Kaplan to defend a certain view of indirect discourse: a report only has to take into account the content, and not the character, of the original utterance (or thought). In particular, indexicals in indirect discourse do not pick out their referents from the context in which the original utterance took place, but from the context of the report: propositional attitude verbs, as stated above, operate on content. Another way to put it is this: since contexts are parameters used to generate content, and content amounts to an intension, Kaplan’s ban on monsters amounts to a prohibition of hyperintensional operators (Westerståhl 2012; Rabern 2012, 2013, 2014; McCullagh 2017; Rabern and Ball 2019). In Kaplan’s own words, “*all operators that can be given an English reading are at most intensional*” (Kaplan 1989a: 502).

3 The Data

Now consider the two following stories:

Situation A: Agnes, a famous politician, sees a recording of one of her speeches. Amazed by her own ideals and rhetoric skills, she thinks to herself: “I will be elected”.

Situation B: Agnes, a thoroughly intoxicated famous politician, sees a woman giving a speech on TV. Amazed by her ideals and rhetoric skills, Agnes thinks to herself: “She will be elected”. Unbeknownst to her, that woman is Agnes herself.

In either situation, it is correct to report Agnes’s belief as (3) below:

(3) Agnes believes that she will be elected

This fact led Kaplan (1989a) to conclude that the self-conscious, indexical nature of the original attribution in situation A is systematically lost when such attribution is reported: in his own terminology, belief reports do not take into account the character of the original sentence. Thus, a belief report of the form “*x* attitude verb that *p*” is true in *c*, *w*, *t* iff the *content* of *p* in *c* is the same as the *content* of *x*’s original utterance. This claim has been largely disputed by later research, most notably on the grounds of the behavior of personal pronouns in several languages (Schlenker 2003; Anand and Nevins 2004) and the functioning of English constructions involving PRO. Let us consider first a counterexample of the latter type, first noticed by Chierchia (1989):

(3a) Agnes expects to be elected

Intuitively, (3a) is true in situation A, but not in situation B. It seems, then, that the truth-conditional import of PRO constructions in attitude reports is sensitive to the self-conscious nature of the original utterances. In addition to this, Chierchia

also noted that sentences like (3) are actually ambiguous between a *de se* and a *de re* reading, and suggested that this difference could be traced to two different LF's. I will come back to this later.

Let us now move to the second kind of counterevidence. Schlenker (2002, 2003) and Anand and Nevins (2004) present relevant data from a variety of languages in which indexicals display a behavior that departs from what LD predicted. For example, the Amharic equivalent of sentence (3b) has two readings: one that is identical to its English reading, and another in which Agnes says that Agnes will be elected:

(3b) Agnes says that *I* will be elected

I won't go much into discussing this kind of pronouns. The data provided by Schlenker does not seem to entail that the Amharic equivalent (3b) is sensitive to the original utterance's character, i.e., it seems that the shifted reading of (3b) is true in both situations above – therefore its relevance to this essay is not much greater than that of the ambiguous (3). Schlenker, however, also provides evidence that some English temporal indexicals – namely, “ago” and “in” – display a behavior analogous to the Amharic first-person pronoun. Even though the indexical nature of such expressions has been called into question (see Anand and Nevins 2004: 34–5), I will be arguing that they can also be understood as cases of binding – an idea already argued for by von Stechow (2002).

In addition to the English and Amharic versions of “I”, Schlenker's criticism of Kaplan also relies on a third type of pronoun, logophoric pronouns. The phenomenon of logophoricity was first described by Clements (1975), and it is just the natural-language version of what Castañeda (1966, 1967) called *quasi-indication*: quasi-indexicals are indexical expressions used in indirect speech to mark the place that a first-person pronoun would have occupied in direct speech. Since such pronouns do not have their own lexical entry in English, Castañeda suggested signaling with an asterisk the occurrences of English demonstratives treated as quasi-indexicals. Thus, a sentence like (3c) would be true in situation A above, but not in situation B:

(3c) Agnes believes that *she*_{*} will be elected

Another way to put it is that (3c) is a disambiguated version of (3). Many languages – most of them belonging to the Niger-Congo family – have lexical entries for pronouns like “*she*_{*}”. Since these are unambiguously *de se*, and therefore sensitive to the character of the original utterance, they may be presented as counterevidence to Kaplan's ideas.

This sort of counterevidence is not limited to uses of the first-person pronoun. Consider the following example:

Situation C: Agnes, a famous politician, has been involved in a case of corruption. Mary, a citizen, sees her walking by the street and tells her “You should resign”.

Situation D: Agnes, a famous politician, has been involved in a case of corruption. Mary sees a bystander and tells her “Agnes should resign”. Unbeknownst to Mary, that bystander is Agnes herself.

This case is analogous to the previous one: sentence (4) is true in both situations, but (4a) is true just in situation C:

(4) Mary told Agnes that she should resign

(4a) Mary told Agnes to resign

Just as essentially first-person indexical attitudes are called “*de se*”, attitudes whose expression involves an essential second-person indexical are usually called “*de te*”. Finally, Schlenker (2003) also mentions a final sort of counterevidence, namely cases of shifted temporal indexicals, i.e., cases in which the report is sensitive to the time of the context of the reported utterance. Consider situations E and F (adapted from Schlenker 2003: 66) and the contrast between (5) and (5a):

Situation E: Agnes, perfectly aware it is Wednesday, reads on a newspaper that it rained last Monday. She tells Rita: ‘It rained the day before yesterday.’

Situation F: Agnes, drunk again, thinks that it is Thursday, when in fact it is Wednesday. She tells Rita: ‘I read on the newspaper that it rained on Monday’.

(5) Agnes said that it had rained two days earlier/before

(5a) Agnes said that it had rained two days ago

Again, (5) is true on both situations, but this is not the case of (5a), which is true only in situation E. What this means is that “ago” is a temporal logophor, i.e., that it is sensitive to the time of the original context – i.e., it is true only if the original utterance could have been made using an indexical like “the day before yesterday”.³

More generally, it is possible to describe the problem of *de se* attitudes as a problem for what Perry (1979: 5-8) calls *the doctrine of propositions*, which consists on the conjunction of the following three claims (I shall use Ninan’s (2016: 92) recent formulation):

TWO-PLACE RELATION: an attitude relation such as believing or desiring is a two-place relation between an agent and a proposition.

FREGE’S CONSTRAINT: if an agent has an attitude that she can express by means of an utterance that has proposition *p* as its content, without having an attitude that she would express by saying that *q*, then *p* and *q* are not the same proposition (i.e., if a rational agent can accept “Hesperus is Phosphorus” without accepting “Venus is Phosphorus”, then these two sentences do not express the same proposition).

³Schlenker also argues that the preposition “in” as in “Rita said that it would rain in two days” can to a great extent be interpreted analogously, but the indexical nature of this word has been called into question because of its dependence on time operators (Anand and Nevins 2004: 34–5).

ABSOLUTENESS: the contents of attitudes are absolute, i.e., they do not vary in truth-value across agents.⁴

There is also an important *rationale* behind ABSOLUTENESS, which is the idea that two people agree on something in virtue of sharing a belief (or any other attitude) with the same content (i.e., an attitude towards the same proposition). Ninan (2016: 100) dubs this idea AGREEMENT:

AGREEMENT: let A_{xp} be a token attitude of an individual x towards a proposition p . Then, for any attitude A and individuals x and y , x and y agree on something in virtue of holding an attitude A_{xp} and A_{yq} iff $p=q$.

The idea behind all these requirements is the following: if (3) is true in situation B but (3a) and (3c) are not, then the embedded sentences that each of them contains do not express the same thing (by FREGE'S CONSTRAINT). But the only way to achieve this result is to either say that the proposition expressed by them depends on some sort of agent parameter (which involves rejecting ABSOLUTENESS) or that "expect" and "believe" operate on something more fine-grained than propositions (which entails a rejection of TWO-PLACE). And if we want the object of Agnes's belief to be just a traditional proposition whose truth value does not vary from agent to agent (the absolute proposition *that Agnes will be elected*), it is unclear why (3c) and (3a) are not true in both situations: what she expects and believes in both cases is the same proposition⁵ – i.e., it is unclear why FREGE'S CONSTRAINT would not hold. Notice that none of these claims presuppose anything about the nature of propositions: they might be Fregean senses, structured entities or intensions – and, even though I will be stating my claims by means of the latter, it is not hard to adapt what I will say to other views of the nature of propositions.

Several proposals have been put forward for handling this problem. Perry (1979) modified TWO-PLACE and argued that beliefs were two-place relations between agents and *belief states*,⁶ which are, roughly, a mental counterpart of Kaplan's character. This involves rejecting ABSOLUTENESS, since the same belief state may determine two distinct propositions with different truth-conditions, depending

⁴Ninan and Perry also want propositions to be invariant across times, but this requirement is incompatible with relativizing the extension of a proposition to a world and a time. Since I am stating my claims within Kaplan's well-known system, I shall drop this requirement, but all the claims made here can easily be adapted to a more conventional framework that relativizes extensions to just possible worlds.

⁵Notice that the problem here cannot be traced to some difference between "believe" and "expect", for English seems to be an exception in not licensing sentences of the form "*a* believes to be *P*", which are common, for instance, in Romance languages (Chierchia 1989: 20). Besides, appealing to such a solution would leave the problem unsolved when it comes to verbs like "desire", which admit both *that*-clauses and infinitive clauses.

⁶It is also usual to read Perry's proposal as a rejection of TWO-PLACE, i.e., as a proposal to understand belief as a three-place relation between an agent, a belief state and a proposition. However, Stalnaker (1981: 149, footnote 11) made clear that this is not the way Perry intended to be understood. Besides, Perry has recently changed his mind (see de Ponte et al. 2019, this volume).

on who holds that state. In reporting such state, it is enough to report the proposition associated with it. Lewis' proposal (Lewis 1979) was to regard *de se* attitudes as self-ascriptions of properties, and, under the supposition that to each proposition it corresponds a property, argued in favor of abandoning propositions; this proposal also amounts to rejecting ABSOLUTENESS, since properties may be true of some individuals but not of others, but here too a version of TWO-PLACE may be retained, since attitudes are still two-place relations between individuals and properties. Chierchia (1989) partially adopted Lewis' idea, and claimed that distinct linguistic constructions were involved in each case: according to him, *de re* and *de se* attitudes are, respectively, attitudes towards propositions and towards properties. Another strategy, put forward by Anand and Nevins (2004) and further developed by Ninan (2010),⁷ is to enrich the *circumstance* parameter with an agent argument. This idea is also Lewisian, since it corresponds to the centered worlds approach that he sketches, and it amounts, again, to rejecting ABSOLUTENESS – the truth value of a proposition is relative to a world and an agent.

My proposal is a bit different. Like Chierchia, I intend to show that it is possible to associate *de se* and *de re* readings of an attitude report with distinct linguistic constructions. But, unlike him, I shall argue that it is possible to achieve a similar result with just the traditional notion of proposition. Namely, I will be claiming that *de se* readings of these sentences can be traced to the behavior of variables at the level of LF. My account requires no context-shifting (unless we follow Kaplan's suggestion to include the assignment function within the context), but it is still monstrous, since it requires a different, more interesting kind of monsters, one dubbed "content monsters" (Rabern and Ball 2019) (I shall define this notion in Sect. 5). Since my point is linguistic rather than epistemological or metaphysical, let me restate the above doctrine as follows:

LING TWO-PLACE: attitude verbs are of type $\langle\langle s, t \rangle, \langle e, \langle s, t \rangle \rangle\rangle$, and they do not undergo type-shifting.

LING FREGE'S CONSTRAINT: if any two sentences of the form "*a* attitude-verb *p*" and "*a* attitude-verb *q*" fail to have the same truth-value, then *p* and *q* do not say the same thing.

LING ABSOLUTENESS: the truth-value of a proposition is not dependent on an agent nor an audience, i.e., circumstances consist of just possible worlds and times (see footnote 4), and do not incorporate individuals of any sort.

LING TWO-PLACE is formulated as in Heim and Fintel's (2011) handbook, which builds on and expands Heim and Kratzer's (1998) standard terminology, which I shall be using throughout this paper. It amounts to saying that the contribution of an attitude verb to what is said is a function from propositions (type $\langle s, t \rangle$) to functions of type $\langle e, \langle s, t \rangle \rangle$, i.e., functions from individuals to

⁷Ninan takes this idea much further. He shows that agent-centered approaches are not fine-grained enough for accounting for the dynamics of conversation, and that at least an addressee parameter is required in order to keep track of what has been said.

propositions. Propositions range over circumstances (type s) and return truth-values (type t). This means that a verb like “believe” picks out a proposition p and returns the property of believing that p , which in turn picks out an individual a (type e) and returns the proposition that a believes that p , which is also of type $\langle s, t \rangle$. We should also introduce some verbs like “tell”, which relate two agents to a proposition – they are of the rather complicated type $\langle e, \langle \langle s, t \rangle, \langle e, \langle s, t \rangle \rangle \rangle$. Chierchia’s proposal amounts to rejecting LING TWO-PLACE as I formulated it, as in his theory intensional verbs undergo type-shifting: sometimes they combine with propositions and sometimes with properties.

LING FREGE’S CONSTRAINT is formulated here in a vague way, since I have employed the locution “say the same thing” instead of “express the same proposition”; the reasons for this will be made clear below. And LING ABSOLUTENESS is clear enough by itself. It is also worth saying a word about AGREEMENT: even if it may not be considered as a requirement for the doctrine of propositions, it is clear that it is a source of data that any theory dealing with the phenomenon of the *de se* should be able to handle. In particular, linguistic phenomena having to do with ellipsis, anaphora and indirect speech, which require a single content to be attributed to several agents, can be used as evidence that a certain theory is on the right track. In particular, consider the following paradigmatic example, due to Perry (1977):

(6) Heimson believes that he is Hume

(6a) Hume believes that he is Hume

Hume and mad Heimson agree on their beliefs. What this means is that we want to be able to say things like the following:

(6b) Heimson believes that he is Hume, and so does Hume

(6c) Both Heimson and Hume believe that they are Hume

The only way to make sense of the meaning of sentences like (6b) and (6c) is by attributing to both Heimson and Hume the same content, even if the former believes falsely and the latter does so truly. Any minimally plausible theory about the *de se* should be able to derive the right truth-conditions for examples of this kind.

So my aim in this paper is to show that the traditional resources of semantics are enough for making all these claims compatible, even if, as I explain at the end of Sect. 5.3, my solution entails a rejection of ABSOLUTENESS at the epistemological level.

Now, a *caveat*: it has been argued that the problem that *de se* attitudes raise for this doctrine is not greater than the classical Frege cases (see for instance Cappelen and Dever 2013), a conclusion that even Ninan, himself a *de se* exceptionalist, seems to embrace.⁸ After all, it is possible to restate a problem analogous to Agnes’s that

⁸He is a *de se* exceptionalist in that he thinks that *de se* attitudes do raise a special problem for a traditional view of propositions, but only if we co-join Perry’s three requirements with

employs just proper names, with no indexicals involved: someone may believe that George Bush is running for the election without believing that Dubya is doing so, and this is a problem for the doctrine above as much as the problem of the *de se* is. Of course, it would in turn raise analogous, well-known problems in attitude reports. It is after all unsurprising that a problem for the general doctrine of propositions is reproduced when we are trying to report those very propositions.

Be it as it may, these concerns are not relevant for us now, since they are orthogonal to the purposes of this essay: if the problem of *de se* attitudes is just a special case of the problem of co-referential terms in intensional contexts, then so much better for semantics, since it has the resources for partially solving this problem, as I shall argue in the following two sections. In Sect. 4 I show that all the sentences involved can be associated to two distinct LF's, each corresponding to a *de se* and a *de re* reading. Section 5 is the core of this paper, since there I discuss key cases of the problem of *de se* attitudes and argue that this analysis, which is traditional and well-established, suffices for dealing with the problem as a whole.

4 Anaphora, Direct Reference and *De Se* Readings

What I intend to show in this paper is that the *de se* readings of sentences can all be associated with a certain kind of LF, namely one in which the variables contributed by the pronouns are bound.

Let us begin with an analysis of (3). As Chierchia noted, this sentence is actually ambiguous between a reading that is true both in situation A and B and another in which it is true just in situation A, the latter reading being roughly an English equivalent of (3c). Chierchia's proposal was to partially adopt a Lewisian framework for *de se* attitudes in which these are seen as self-ascriptions of properties. Thus, Chierchia would claim that a sentence like (3) is ambiguous between (3d) and (3e):

(3d) believe (Agnes, λx . x will be elected)

(3e) λx . believe (x, x will be elected) (Agnes)

In (3d), "believe" is construed as a relation between an agent and a property, corresponding to the *de se* reading, whereas in (3e) it is construed as a relation between an agent and a proposition. This kind of approach has been criticized for its binary character, and indeed I will be pursuing here an approach that, I expect, will be able to handle every case with just propositions.

AGREEMENT and an additional requirement which he dubs "EXPLANATION": the idea that action is explained in terms of beliefs and desires (Ninan 2016: 99–109). As I see it, this is a way to shift the focus from philosophy of language to epistemology and philosophy of mind and action.

First, notice that there are two ways in which “she” in (3) can be anaphoric on “Agnes”: either semantic binding or simple co-reference.⁹ The technical details about this issue are rather well-known in the literature of semantics, but I shall restate them here. Also, I will be omitting relativization to contexts and circumstances for simplicity’s sake, so pretend for the moment that we are working simply with an extensional framework. In Heim and Kratzer’s (1998) standard terminology, a sentence’s main verb phrase is an unsaturated entity: an expression of type $\langle e, t \rangle$, i.e., a function ranging over objects and yielding truth values. Noun phrases, on the other hand, are normally understood as type e : they yield an object, and thus the composition of noun and verb phrases yields truth values, as desired:

(7) John hates Paul

However, quantified noun phrases pose a problem for this simple idea, since they do not yield any object, and therefore a sentence like (7a) would fail to receive a truth value, for “hates Paul” cannot take “everyone” as its object:

(7a) Everyone hates Paul

The solution, due to Montague (1973), lies in considering quantified phrases as expressions of type $\langle \langle e, t \rangle, t \rangle$, i.e., as functions ranging over properties and returning truth values. Thus, whereas the verb phrase of (7) would take “John” as its argument to yield a truth value, the roles are reversed in (7a), in which the verb phrase “hates Paul” would be the argument of the noun phrase “Everyone”. To repeat, “everyone”, of type $\langle \langle e, t \rangle, t \rangle$, takes as its argument a function of type $\langle e, t \rangle$ such as “hates Paul”, and returns thereafter a truth value. This idea, however, poses yet another problem:

(7b) Paul hates everyone in this room

A transitive verb like “hates” is of type $\langle e, \langle e, t \rangle \rangle$, i.e., it picks an object a and returns the property of hating a . But, again, “everyone in this room” cannot compose with “hates” in this way, since it yields no object. The solution to this is to postulate a kind of movement occurring at the level of LF, which receives the name “Quantifier Raising” (from now on, “QR”). The idea is basically this: a quantified phrase can move to any other place within the sentence, always leaving behind a trace t and a lambda operator that binds the variable contributed by the trace. Thus, (7b) would receive the following LF and truth-conditions:

LF: [*Everyone in this room*] λi . Paul hates t_i
 [[[*Everyone in this room*] λi . Paul hates t_i]]^g = 1 iff

⁹Kaplan (1989b: 572) claims that anaphoric and exophoric (deictic) uses of demonstratives correspond, respectively, to bound and free uses of variables. However, this analogy has been called into question, since many cases traditionally considered as anaphora are better analyzed as cases of free variables that just happen to be coreferential with some previously appeared item (Heim and Kratzer 1998: 240).

$[[\text{Everyone in this room}]]^g ([[\lambda i. \text{Paul hates } t_i]])^g = 1$ iff
 $[\lambda P. \text{Every } x (x \text{ is in this room} \rightarrow P(x)) (\lambda y. \text{Paul hates } y)]$

Now, the fact that proper names can be combined with quantified phrases, as in (7c) below, provides evidence that they can undergo *type-shifting*, i.e., that they can be interpreted as expressions of type $\langle\langle e, t \rangle, t \rangle$:

(7c) John and everyone in this room hate Paul

LF of (7c): $[\text{John and everyone in this room}] \lambda i. t_i \text{ hates Paul}$

In other words, proper names can be treated as quantifiers. Now, semantic binding, either in the case of quantifiers or of proper names, can only be achieved by means of QR, i.e., by means of letting the lambda abstractor bind not just the trace but also whatever pronoun we want to interpret as bound. This means there are two different ways to achieve an anaphoric reading of (7d) below: one in which “Paul” is QR’d and binds the variable contributed by “his”, and another one in which “his” is left as a free variable that just happens to co-refer with “Paul”.¹⁰ E.g., let g be an assignment such that $g(i) = \text{Paul}$:

(7d) Paul_i hates his_i father

Co-reference: $[[\text{Paul hates his}_i \text{ father}]]^g = 1$ iff
 $[[\text{hates his}_i \text{ father}]]^g ([[\text{Paul}]])^g = 1$ iff
 $\lambda x. x \text{ hates Paul's father (Paul)}$ iff
 Paul hates Paul ‘s father

Binding: $[[[\text{Paul}] \lambda i t_i \text{ hates his}_i \text{ father}]]^g = 1$ iff
 $[[\text{Paul}]]^g ([[\lambda i t_i \text{ hates his}_i \text{ father}]])^g = 1$ iff
 $[\lambda P. P(\text{Paul})] (\lambda x. x \text{ hates } x \text{'s father})$ iff
 $\lambda x. x \text{ hates } x \text{'s father (Paul)}$ iff
 Paul hates Paul’s father

To each of these two possibilities it corresponds a different LF. This kind of syntactic ambiguity is responsible for a number of linguistic phenomena; e.g., bound variable readings and free-variable readings are responsible, respectively, for the so-called “sloppy” and “strict” readings of ellipsis.¹¹ Now, the very same analysis

¹⁰Notice that there is yet another LF available, one in which “John” is raised but the lambda operator happens to bind only “John”’s trace. This LF is analogous to that of the free-variable reading.

¹¹To see why, consider the following sentence:

(7e) John hates his father, and so does Bill

Excluding those cases in which “his” is not anaphoric on “John”, there are two readings available for the second conjunct in (7e): one in which Bill hates John’s father (*strict reading*) and one in which he hates his own father (*sloppy reading*). It is easy to see how each of them stems from copying the content of the VP applied to “John” in, respectively, the co-referential and the bound

can be applied to (3). Let us adopt an intensional system again: let C be the set of circumstances (i.e., the Cartesian product of the set W of possible worlds and the set T of times) and let Π be the set of propositions (i.e., the set of functions $W \times T \rightarrow \{1, 0\}$). Also, let h be an assignment such that $h(i) = \text{Agnes}$:

Free Variable:

LF: *Agnes believes that she_i will be elected*

$[[\text{Agnes believes that she}_i \text{ will be elected}]]^{c, h, w, t} = 1$ iff

$[[\text{believes that she}_i \text{ will be elected}]]^{c, h, w, t} ([[\text{Agnes}]])^{c, h, w, t} = 1$ iff

$[[\text{believes}]]^{c, h, w, t} (\lambda \langle w_1, t_1 \rangle \in C. [[\text{she}_i \text{ will be elected}]]^{c, h, w_1, t_1})$
 $([[\text{Agnes}]])^{c, h, w, t} = 1$ iff

$[[\text{believes}]]^{c, h, w, t} (\lambda \langle w_1, t_1 \rangle \in C. h(i) \text{ will be elected in } \langle w_1, t_1 \rangle) (\text{Agnes}) = 1$
 iff

$[\lambda \pi \in \Pi. [\lambda x. \text{for all circumstances } \langle w_2, t_2 \rangle \text{ compatible with what } x \text{ believes}$
 $\text{in } \langle w, t \rangle, \pi(\langle w_2, t_2 \rangle) = 1 (\lambda \langle w_1, t_1 \rangle \in C. h(i) \text{ will be elected in } \langle w_1, t_1 \rangle)$
 $(\text{Agnes}) = 1$ iff

In all circumstances $\langle w_2, t_2 \rangle$ compatible with what Agnes believes in $\langle w, t \rangle$,
 $h(i)$ will be elected in $\langle w_2, t_2 \rangle$

Bound Variable:

LF: *[Agnes] $\lambda i. t_i$ believes that she_i will be elected*

$[[[\text{Agnes}] \lambda i. t_i \text{ believes that she}_i \text{ will be elected}]]^{c, h, w, t} = 1$ iff

$[[[\text{Agnes}]]^{c, h, w, t} ([[\lambda i. t_i \text{ believes that she}_i \text{ will be elected}]])^{c, h, w, t})$
 $= 1$ iff

$[[[\text{Agnes}]]^{c, h, w, t} (\lambda x. [[\text{believes}]]^{c, h(x/i), w, t} (\lambda \langle w_1, t_1 \rangle \in C. [[\text{she}_i \text{ will be}$
 $\text{elected}]]^{c, h(x/i), w_1, t_1}) ([[t_i]]^{c, h(x/i), w_1, t_1})) = 1$ iff

$[\lambda P. P(\text{Agnes}) \text{ in } \langle w, t \rangle] (\lambda x. [\lambda \pi \in \Pi. [\lambda y. \pi(\langle w_1, t_1 \rangle) = 1 \text{ for all circum-}$
 $\text{stances } \langle w_1, t_1 \rangle \text{ compatible with what } y \text{ believes in } \langle w, t \rangle]] (\lambda \langle w_2, t_2 \rangle \in$
 $I. x \text{ will be elected in } \langle w_2, t_2 \rangle) (x) = 1$ iff

$[\lambda P. P(\text{Agnes}) \text{ in } \langle w, t \rangle] (\lambda x. x \text{ will be elected in all circumstances } \langle w_1, t_1 \rangle$
 $\text{compatible with what } x \text{ believes in } \langle w, t \rangle)$

Notice that there is a notable difference between these two LF's: if "she" happens to be bound, it must always be anaphoric on "Agnes", but this needn't be the case if it is free. In particular, if "she" is free its semantic value will depend on the assignment function, which means that it can refer to any salient female. Notice also that, if "she" is free but does not refer to Agnes (say, because the utterer is pointing at some other person), then the belief reported by (3) is always *de re*.

readings. On the sloppy reading, "Bill" is QR'd and yields an abstractor to the logical that binds the pronoun in "hates x 's father", thus returning a property that has nothing to do with John (namely, the property of being an x that hates x 's father). The strict reading, on the other hand, results simply from copying the content of the VP applied to "John" on the directly referential reading of the first conjunct under that very same assignment, i.e., it attributes to Bill a property that does have to do with "John", namely hating John's father.

My proposal is straightforward: it is possible to assign distinct LF's to the *de se* and to all the *de re* readings of (3): if the variable is left free, it follows that it can refer to any individual depending on the assignment, in which case it seems like a proper analysis for all the possible *de re* readings of (3), including the case in which $h(i) = \text{Agnes}$. Thus, all the *de re* readings can be traced to a single LF. If, on the other hand, "she" is bound, there is only a reading available, and thus it can be happily assumed that it corresponds to the only possible *de se* reading. As we shall see in the next section, this difference, often overlooked, can play the role attributed to properties or enriched propositions in other accounts of the *de se* – namely, because lambda binders are monsters (Rabern 2013).

My analysis departs from Chierchia's in that "believes" is construed, on either reading, as a relation between a subject and a proposition. The main difference lies in the treatment of the subject itself: on the *de re* reading, "Agnes" works as a proper name, yielding an object, whereas on the *de se* reading it is treated as a quantifier and raised, leaving a trace behind and introducing a lambda abstractor to the LF.

Now, this very same analysis can easily be extended to PRO and logophoricity. The latter case is rather easy: recall that logophoric pronouns are just natural language counterparts of "she_{*}" in (3c). A sentence like (3c) has the same truth-conditions, and the same superficial structure, as the *de se* reading of (3); therefore it seems appropriate to say that both should share an LF, i.e., that "she_{*}" can only be read as a bound variable. This fact is further supported by the idea that, as Castañeda (1966, 1967) speculated and the empirical studies by Clements (1975) made clear, logophors are always anaphoric.

On the other hand, PRO is a null pronominal form, which means that it contributes a variable to LF. Given that PRO is actually a logophor, my argument above holds for this case too: the variable contributed by PRO has only a bound reading. This yields the following LF and truth-conditions for (3a):

LF: [Agnes] $\lambda i. t_i$ expects PRO_i be elected
 $[[[\text{Agnes}] \lambda i. t_i \text{ expects } \text{PRO}_i \text{ be elected}]]^{c, g, w, t} = 1$ iff
 $[\lambda P. P(\text{Agnes}) \text{ in } \langle w, t \rangle] (\lambda x. x \text{ is elected in all circumstances } \langle w_1, t_1 \rangle \text{ compatible with what } x \text{ expects in } \langle w, t \rangle)$

It is easy to check that this form is analogous to the *de se* reading of (3), which means that it is possible to associate a *de se* and *de re* readings of a sentence with different linguistic patterns: namely, bound variables trigger the former, and free variables are associated with the latter.

Now, the LF's that I have provided for *de se* readings of sentences are strictly analogous to Chierchia's (3e), a form that Chierchia explicitly reserved for *de re* readings. He raises a number of objections against the idea that a propositional analysis can be provided for *de se* readings, and argues that an LF like (3e), in which "believe" relates an agent to a proposition, will never yield the right results. My aim in the next section is to answer all these challenges: I will be arguing that (3e) is indeed a good analysis of the *de se*. After discussing some of the best-known examples of *de se* attitudes, I shall extend my analysis to sentences such as (4a), (5a) and (6–6a).

5 Monsters and the Objects of Belief

5.1 *The Problems for Propositionalism*

Chierchia's objections against treating propositions as the objects of *de se* attitudes are essentially variants of Lewis's. Consider sentences (6–6a), repeated below:

(6) Heimson believes that he is Hume

(6a) Hume believes that he is Hume

As Lewis claims, Hume and Heimson believe the same thing. However, this result cannot be achieved by means of propositions: “he” in (6) is co-referential with Heimson, whereas in (6a) it refers to Hume; a propositional analysis would yield the result that each of them believes a different proposition, namely that Heimson is Hume and that Hume is Hume, respectively. But, in this case, Heimson and Hume do not believe the same thing, and we cannot make sense of sentences like (6b) and (6c). Indeed, this analysis seems wrong, since it attributes to Heimson a belief that he does not actually hold – that Heimson is Hume. And, even if someone could make sense of the idea that both Heimson and Hume believe the same proposition, she would still have to explain why the former believes falsely and the latter truly (i.e., why ABSOLUTENESS does not hold). Lewis's and Chierchia's solution is, of course, to treat “believe” as a relation between an agent and a property, which yields the right results: both Heimson and Hume self-attribute the property of being Hume.

A similar objection is raised by Chierchia against treating PRO constructions as propositions. In his proposal, PRO is understood as a lambda abstractor, and infinitives and gerunds as properties. Consider his own example:

(8) Pavarotti practiced singing *Rigoletto*

(8a) Domingo practiced it too (what Pavarotti practiced)

If “practice” were understood as relating a subject to a proposition, (8) would receive an analysis as in (8b) below. But the anaphora in (8a) must be resolved by copying whatever object is attributed to Pavarotti, which yields the wrong result, since it relates Domingo to the proposition that Pavarotti sang *Rigoletto*, as in (8c):

(8b) $\lambda x.$ practice (x , x sings *Rigoletto*) (Pavarotti)

(8c) practice (Domingo, Pavarotti sings *Rigoletto*)

In other words, only sloppy readings should be allowed for cases of ellipsis and anaphora involving PRO and infinitives or gerunds. Hence, Chierchia concludes, the contribution of constructions of the form PRO+non-finite clause should be a property instead of a proposition, namely one like (8d):

(8d) $\lambda x.$ x sings *Rigoletto*

If the anaphora in (8a) is resolved as (8d) instead of (8b), we get the right results. But, as we shall see now, an analysis in propositional terms is not ruled out by this kind of counterexamples. In the next subsection I will argue that any propositional account that relates a subject with a proposition containing a bound variable can – and indeed will – yield the right results for sentences like these.

5.2 *Monsters, Pronouns and Compositionality*

Recall the above discussion on monsters. The contribution of “I” to what is said varies depending on whether it is found isolated or embedded under the monstrous operator M . Imagine that (2) and (2a) are uttered in a context having Angela Merkel as its agent:

(2) I am tired

(2a) M (I am tired)

The content contributed by “I” in (2) is Angela Merkel. However, what is said by (2a) is not a function of this content: monsters are hyperintensional operators, and as such they prevent the content of indexicals from composing into what is said by a sentence, i.e., they shift the contribution made by the indexicals under their scope. Indeed, it has been argued that the three characterizations of a monster provided by Kaplan (Predelli 2014) are just versions of a single notion of monster: an operator that shifts the contribution that an expression makes to assertoric content (Rabern 2014; McCullagh 2017). The traditional understanding of a monster as an operator on context boils down to a hyperintensional operator, since contexts are tuples of parameters that generate *content*, and therefore a shift in context may prevent sentences from being compositional at this level of meaning. This also seems to be what Kaplan has in mind when he talks about “operators on character”, since character is a function that generates content out of context.

A series of recent papers (Rabern 2012, 2013; Rabern and Ball 2019) have put forward a seemingly problematic claim about Kaplan’s LD: namely, that quantifiers, and indeed all variable binders, are monsters. The argument goes as follows: variable binders shift the assignment function, which is *a parameter that generates content*. In Kaplan, the content of the variables is given by the assignment function: if α is a variable, then $\lambda w. \lambda t. [[\alpha]]^{c, g, w, t}$ is a constant function from circumstances to $g(\alpha)$, just like $\lambda w. \lambda t. [[I]]^{c, g, w, t}$ is a constant function from circumstances to c_A . Content, then, is dependent not only on context, but also on assignment. Therefore operators like Tarskian quantifiers, which shift the assignment function, have the same effect as operators that shift context: both kinds of operators prevent sentences from being compositional at the level of content – they are hyperintensional. Let us see an example; consider (9) and (9a):

(9) Px

(9a) $\exists x (Px)$

In Kaplan's system, the content of (9) under assignment g is the function that assigns truth to all the circumstances in which $g(x)$ is P. But $g(x)$ does not compose into what is said by (9a), whose truth conditions are as follows:

$[[\exists x (Px)]]^{c, g, w, t} = 1$ iff there is some a such that $[[Px]]^{c, g(a/x), w, t} = 1$, where $g(a/x)$ is just like g except that it assigns a to x .

To put it otherwise, (9) and (9b) have the same content under an assignment g such that $g(x) = g(y)$, but they embed differently under the existential quantifier: (9a) and (9c) do not mean the same, which means that the existential quantifier does not operate on content:

(9b) Py

(9c) $\exists x (Py)$

In sum: variable binders shift the assignment function, and this function is a parameter that generates content. Therefore quantifiers and other variable binders are hyperintensional operators, and have the same effect as operators on context. This has led Rabern and Ball to claim that variable binders are indeed monsters. Indeed, Kaplan even goes on to suggest that we include the assignment function among the parameters of context (Kaplan 1989b: 593); if we do this, then variable binders are operators on context – and monsters in the traditional sense, after all. But Rabern and Ball's point is even stronger: even if we stick to the formalism of *Demonstratives* and place assignment functions outside context, variable binders still deserve the name “monsters”, since they operate on characters just as context-shifters do. Operators on context are interesting only insofar as they are *content monsters*: devices that shift a parameter on which assertoric content depends. And this is precisely what assignment-shifters do.

Now, it may be argued that Kaplan may get rid of this inessential inaccuracy of LD by doing away with the old Tarskian quantifiers and adopting generalized quantifiers, as it is commonplace (and indeed necessary) in semantics. But this will not do either since, as we have seen above, semantic binding can only be achieved by means of lambda abstractors, which are, in turn, assignment shifters. Namely,

$[[\lambda i. \Phi]]^g = \lambda x. [[\Phi]]^{g(x/i)}$ (Heim and Kratzer 1998: 196; Rabern 2013: 399).

Having this in mind, my response to the objections by Lewis and Chierchia is straightforward, as we shall see in the next subsection.

5.3 *Believing the Same*

The objection by Chierchia and Lewis was that a propositional analysis of (6) and (6a) could never yield the result that Heimson and Hume are attributed the same belief. But now consider my analysis of such sentences:

LF of (6): $[Heimson] \lambda i. t_i$ believes that he_i is Hume

LF of (6a): [Hume] $\lambda i. t_i$ believes that he_i is Hume

Focus your attention on the embedded sentence, i.e., “ he_i is Hume”. This sentence is supposed to represent the content of the belief. Chierchia’s and Lewis’s complaint amounts to saying that, if “he” is co-referential with the subject of the main sentence, then they cannot believe the same thing. But actually this pronoun is not co-referential with anyone, since it is not referential at all: it is a bound variable. Bound variables are not referential devices. In particular, the content of the embedded sentence is about neither Hume nor Heimson; the binding has the effect that $g(i)$, whoever she is, becomes irrelevant for assessing the content of (6) and (6a). Therefore Heimson does not appear anywhere within what is said by the embedded clause.

My analysis here is completely parallel to the standard analysis of strict and sloppy resolutions of ellipsis (see footnote 11). Heimson and Hume believe the same thing if the content of Heimson’s belief is the same as the content of Hume’s. Attributing to Heimson the belief in the same proposition as Hume yields the right results if the variable is bound. Calculate as follows (I will be omitting relativization to contexts for simplicity’s sake):

The content of (6) is as follows:

$$\begin{aligned} & \lambda w. \lambda t. [[[\text{Heimson}] \lambda i. t_i \text{ believes that } he_i \text{ is Hume}]]^{g, w, t} = \\ & \lambda w. \lambda t. [[\text{Heimson}]]^{g, w, t} (\lambda w. \lambda t. [[\lambda i. t_i \text{ believes that } he_i \text{ is Hume}]]^{g, w, t}) = \\ & \lambda w. \lambda t. [[\text{Heimson}]]^{g, w, t} (\lambda w. \lambda t. \lambda x. [[t_i \text{ believes that } he_i \text{ is Hume}]]^{g(x/i), w, t}) = \\ & \lambda w. \lambda t. [[\text{Heimson}]]^{g, w, t} (\lambda w. \lambda t. \lambda x. [[\text{believes}]]^{g(x/i), w, t} (\lambda w. \lambda t. [[he_i \text{ is Hume}]]^{g(x/i), w, t}) \\ & \quad (\lambda w. \lambda t. [[t_i]]^{g(x/i), w, t})) \end{aligned}$$

It is clear from this that the content of the proposition believed by Heimson is not about any individual at all. Namely, no element of the domain enters into the content of such belief. In terms of structured propositions, the content of “Heimson believes that he is Hume” could be represented as follows:

(being a property of Heimson, believing that the holder of the belief is Hume)

Now, (6a) receives exactly the same analysis:

$$\begin{aligned} & \lambda w. \lambda t. [[[\text{Hume}] \lambda i. t_i \text{ believes that } he_i \text{ is Hume}]]^{g, w, t} = \\ & \lambda w. \lambda t. [[\text{Hume}]]^{g, w, t} (\lambda w. \lambda t. [[\lambda i. t_i \text{ believes that } he_i \text{ is Hume}]]^{g, w, t}) = \\ & \lambda w. \lambda t. [[\text{Hume}]]^{g, w, t} (\lambda w. \lambda t. \lambda x. [[t_i \text{ believes that } he_i \text{ is Hume}]]^{g(x/i), w, t}) = \\ & \lambda w. \lambda t. [[\text{Hume}]]^{g, w, t} (\lambda w. \lambda t. \lambda x. [[\text{believes}]]^{g(x/i), w, t} (\lambda w. \lambda t. [[he_i \text{ is Hume}]]^{g(x/i), w, t}) \\ & \quad \lambda t. [[t_i]]^{g(x/i), w, t})) \end{aligned}$$

Needless to say, what is believed by Hume is the same as what Heimson believes. In terms of structured propositions, it may be represented as follows:

(being a property of Hume, believing that the holder of the belief is Hume)

It is straightforward that this analysis yields the right resolution of the ellipsis in (6b). They believe the same even if each of their beliefs is fulfilled by a different set of possible worlds. This is so because what is attributed to each of them is a hyperintension, namely a function from assignments to propositions: the abstractor

in LF does the job of “shifting the content” believed by each of them. This, however, does not mean that the verb “believe” cannot be treated as a relation between an agent and a proposition: what this means is simply that this proposition does not compose into what is said by the whole sentence. This amounts to saying that LING TWO-PLACE can be retained, since “believe” is still of the appropriate type, but we needn’t abandon LING FREGE’S CONSTRAINT either: the embedded clause of (6) does not say the same thing as the embedded clause of (6d):

(6d) Heimson believes that Heimson is Hume

Since these two embedded clauses do not share a content, Heimson may be considered rational. However, the embedded clause does say the same thing in (6) and (6a), which means that we can perfectly say that Heimson and Hume *agree* in their beliefs – even if the thing believed is not a proposition. That is why I formulated LING FREGE’S CONSTRAINT in a loose way, by using the vague notion of “saying the same thing” instead of “expressing the same proposition”. Now this principle can be made more precise by stating that if any two sentences of the form “*a* attitude-verb *p*” and “*a* attitude-verb *q*” fail to have the same truth-value, then *p* and *q* do not express the same hyperintension. Of course, there is no need to enrich the circumstance parameter in order to achieve these results, so we can also retain LING ABSOLUTENESS, which means that, if the claims made here are right, it is possible to vindicate a traditional approach to the phenomenon of *de se* attitudes.

Some epistemological and metaphysical remarks are in place. A possible complaint about my proposal may hold that, metaphysically speaking, this cannot be considered a propositional approach, since what is attributed in resolving the ellipsis is not a proposition, at least not in the traditional sense – it is not a function from possible worlds to truth-values but a function from assignments to propositions. This is right, but it needn’t constitute an objection, since my point was not that the object of belief is a proposition, but that it is still possible to treat verbs like “believe” as relating a subject to a proposition. Propositions, in the traditional sense, have been long recognized not to be fine-grained enough for dealing with essential indexicality, the *de se* and other related phenomena,¹² but my point was linguistic rather than metaphysical: all this can be dealt with by means of well-established tools of formal semantics, and there is no need to enrich the notion of proposition nor to assume that some verbs may undergo type-shifting. Of course, what is believed by Hume and Heimson, if I am right, is not a proposition, metaphysically and epistemologically speaking. Indeed, it involves a rejection of epistemological ABSOLUTENESS, since it allows two individuals to be right and wrong towards the same belief at the same time. None of this entails that, linguistically, attitude verbs should not relate agents to propositions: LING TWO-PLACE can be retained.

¹²The classical essays by Lewis (1979) and Perry (1979) are credited for establishing the objections to the doctrine of propositions. See also Stalnaker (1981) for a defense of a propositional approach to these phenomena.

Anyway, my account seems *prima facie* very appropriate from an epistemological point of view: what is attributed to both Heimson and Hume is the belief that the holder of the belief, whoever that person may be, is Hume; what Agnes expects is that whoever has her expectation will be elected. This is still very sketchy, but it looks like an accurate characterization of *de se* attitudes. More precisely, the content of these beliefs may be represented by a function from individuals to propositions, which is something that can be easily argued to correspond to the classical conception of *de se* attitudes as self-locating ascriptions. The fact that (3a) is satisfied by the same set of possible worlds as the *de re* reading of (3) is a contingent matter which depends on the fact that it is Agnes herself that is the holder of the belief, but this is also a consequence of Lewis's and Chierchia's views: even though they employ more fine-grained notions of belief that allows them to discriminate different beliefs, Agnes's *de se* belief is still satisfied by the same set of non-centered possible worlds as its *de re* counterpart. The same degree of fine-grainedness for individuating belief can be achieved, I believe, by means of binding, which makes it unnecessary to treat "believe" as a verb combining with a property instead of a proposition.

Now, all the other problematic examples may be analyzed in a similar fashion. Let us consider first the case of PRO. In my proposal, (8) receives the following LF:

LF of (8): [Pavarotti] $\lambda i. t_i$ practiced PRO_i singing Rigoletto

For similar reasons as above, the variable contributed by PRO is not about Pavarotti or about anyone in particular. Therefore we can treat "practice" as a relation between an agent and a proposition (or a content) and attribute this very same proposition to Domingo. The resolution of this ellipsis will still give the sloppy result, as Chierchia wants it. The same analysis applies to (4a), repeated below:

(4a) Mary told Agnes to resign

LF: [Agnes] $\lambda i. \text{Mary told } t_i \text{ PRO}_i \text{ to resign}$

[[[Agnes] $\lambda i. \text{Mary told } t_i \text{ PRO}_i \text{ to resign}]]^{c, g, w, t} = 1$ iff

[$\lambda P. P(\text{Agnes})$ in $\langle w, t \rangle$] ($\lambda x. x$ resigns in all circumstances $\langle w_1, t_1 \rangle$ compatible with what Mary told x in $\langle w, t \rangle$) = 1

Here, "Agnes" is also QR'd and treated as a quantifier, which gives the result that what Mary tells Agnes has nothing to do with Agnes herself. A feature of PRO in my account is that its antecedent must always be QR'd, and this is precisely what explains that ellipsis and anaphors involving PRO always get a sloppy reading, as in (4b):

(4b) Mary told Agnes to resign. She told that to Rita, too

Mary tells the same thing to both politicians, yet it seems that Mary's request in each case will be fulfilled by different resignations – i.e., different propositions must be the case in order for Mary's request to be fulfilled. Does this call against treating "tell" as a relation between agents and a proposition? No, simply because, even if it relates two agents to a proposition, this proposition does not compose into

what is said by (4a) – the individual contributed by PRO does not enter into the truth-conditions of (4a).

The fact that no individual enters into such content is important, and it can be used to shed light on a point that may indicate that this approach is on the right track. Consider the contrast between (3c) (repeated below) and (3e):

(3c) Both Heimson and Hume believe that they are Hume

(3e) Both Heimson and Hume believe that he is Hume

Sentence (3e) is undoubtedly *de re*: there is no way to achieve a reading of it in which both Heimson and Hume each think of themselves as Hume. This is in sharp contrast with (3c), which is standardly understood *de se* – indeed it would be hard to achieve a *de re* reading of it, as it would involve very weird situations in which Hume is a plurality of individuals. These facts are straightforwardly explained in the present theory if we assume an important – and largely uncontroversial – principle of semantics: Heim’s (2001) principle of Feature Transmission Under Binding (FTUB). To see how it works, consider (10–10c):

(10) The boys each think that they are alone in the room

(10a) We brought our children

(10b) Few men brought their children

(10c) Only I voted for myself

In (10), “they” is morphologically plural, but it is semantically understood as a variable ranging over individuals. (10a) has two readings: one in which each of the persons referred to brought her own children, and another in which they have children in common and bring them. (10b) is ambiguous between a demonstrative use of “their” (say, a use in which the speaker is pointing at someone’s children) and another one in which each of the men (individually) brought his own children. A similar phenomenon occurs in (10c): it may be understood as saying either that c_A is the only x that voted for c_A ’s, or that c_A is the only x that voted for x .¹³ Thus, on the former reading “myself” has its usual functioning, but on the latter it is not first-personal at all, even if it is morphologically so.

Heim’s principle was devised for accounting for this sort of cases. What it states is that a moved phrase (like “Both Heimson and Hume” in the two sentences under discussion) transmits, at the phonological level, all its φ -features (gender, number, person, etc.) to the variables it binds. For example, if a certain variable is bound by a plural quantifier it will be phonetically realized as a plural, even if it is simply a

¹³This difference is far from being idle. Let $c_A = \text{Obama}$. If Trump votes for Obama, the first reading is falsified, since Obama is not the only x that voted for Obama; however, this fact needn’t falsify the second reading. If, on the other hand, Trump votes for Trump, then the second reading is false, since Obama is not the only x that voted for x , but the first one can still be true.

variable ranging over individuals, and any variable bound by a first-person pronoun will be realized as a variant of “I” even if semantically it is not first-personal at all. This is how the ambiguity of (10a), (10b) and (10c) is achieved: if the pronoun is left as a free variable, it retains its φ -features at the semantic level, but if it is bound it is semantically just a bare variable that ranges unrestrictedly over individuals. That is why (10), (10a) and (10b) contain a morphologically plural pronoun that is understood semantically as a singular, and (10c) a morphologically first-person, semantically non-first-personal pronoun.

Now apply this very same analysis to (3c) and (3e). The phrase “both Heimson and Hume” is plural, which means that there is no way that “he” in (3e) be bound by it – otherwise it would be phonologically realized as a plural. On the other hand, the standard reading of (3c) is clearly a case in which the variable contributed by “they” is bound, a fact that explains that it is realized as a plural even if it is, semantically speaking, just a bare variable ranging over individuals. Using the variables of first-order logic, we may represent the two LF’s as follows:

LF for (3e): [Both Heimson and Hume] think that $x^{3rd, singular}$ is Hume

LF for (3c): [Both Heimson and Hume] $\lambda x. x$ think that $x^{3rd, plural}$ are Hume

Here, the superscripts indicate the φ -features of the variable. In the LF of (3c) the variable is bound, and that is why I represented its φ -features erased: what this means is that the variable x can range over all objects individually instead of being restricted to ranging over pluralities, even though phonologically it will be realized as a plural because it is semantically bound by the plural quantifier “both Heimson and Hume”.

In a nutshell: the variable contributed by the pronoun in (3e) can never be understood as bound, whereas the variable in (3c) is clearly bound on its standard reading. On the other hand, (3e) can only be understood *de re*, and (3c) is standardly *de se*. This evidence strongly suggests that *de se* attitudes in natural language are just a special case of variable binding – i.e., it provides evidence that the current theory is on the right track.

Now we have at hand the resources for extending the present analysis to the final sort of counterevidence: cases in which a report is sensitive to the time of the context in which the original utterance took place. Consider sentences (5) and (5a), repeated below:

(5) Agnes said that it had rained two days before

(5a) Agnes said that it had rained two days ago

These examples can also be dealt with as cases of variable binding, similarly to what von Stechow (2002) proposed for the case of “in” in the development of a general theory of tense; the theory I am going to briefly sketch here is essentially

a variant of his,¹⁴ which treats tense and mood extensionally, with time and world variables explicitly represented in LF. I, however, will be keeping an intensional treatment for the modal domain.

A few additional assumptions must be made before we proceed. In the first place, assume that verb tenses are represented as items in LF. I will be using $d_0 \dots d_n$ as variables for times so as to avoid confusion with traces; also, let $g(d)$ be the time picked out by each variable under assignment g . If this is so, nothing prevents us from binding some of these variables, as it is usually done when dealing with tense (tenses, after all, express relations). On the other hand, the theory is extensional, which means that relativization of contents to times is vacuous; therefore we may just think of circumstances as consisting of just a possible world. Now, an important thing is that these variables also have features (present, past and so on), and that these too are subject to FTUB: if they are bound, they are realized as if they had the features of their binders, but they are not semantically interpreted as such. Let us provide some LF's:

LF of "It rains": $d_0^{pres} \text{ rain}$

Here, the superscript to the variable indicates that the tense is present. Present features may be considered semantically as an indexical, i.e., $[[d_n^{pres} p]]^{c, g, w}$ is true at c , g and w iff it is the case that p at c_T and w . This means that, given a context c , an assignment g can only be used for giving a value to variables if it assigns c_T to the variable marked with the present feature. On the other hand, past tenses express relations, which means that they are always relative to a present. This can be accounted for at the presuppositional level: the clause between brackets is to be interpreted as a presupposition, which means that an assignment g can give a value to the variables in the following LF only in case it assigns to " d_1 " an instant of time previous to c_T :

LF of "It rained": $d_1^{past} [d_1 < d_0^{pres}] \text{ rain}$

This LF yields truth when assessed at c , g and w iff and it rains at $g(d_1)$ and w , under the presupposition that $g(d_1)$ occurs before c_T . An LF like the one above entails that a variable for the present time will always be represented at LF. Representing time variables explicitly at LF and therefore making propositions insensitive to the time of assessment (they are just functions of possible worlds) is incompatible with the way I treated "believe" up until now, since it is supposed to combine with temporal propositions (i.e., with functions from worlds and times to truth-values). This can be solved by representing explicitly the kind of quantification over times that these verbs perform:

¹⁴His theory does not discuss just tense: it is also devised for dealing with the phenomenon of the *de se* more generally. He treats attitude verbs as operating on centered propositions (i.e., propositions that incorporate an agent parameter in addition to worlds and times), much like Lewis' proposal sketched above, which means that he rejects LING ABSOLUTENESS. It is nevertheless possible to buy his theory of tense without buying his insights into these issues.

(5b) Agnes believes that it rains

LF of (5b): d_0^{pres} Agnes believes λd_1^{pres} . d_1^{pres} rain

The lambda abstractor is contributed by the verb “believe”, and it binds the present variable contributed by “rain”, which makes the latter’s present feature non-interpretable because of FTUB. Thus Agnes’s thought can be represented as a set of times and worlds in which c_T plays no special role: even though the indicative “rain” would normally pick out c_T , here it works simply as a variable bound by the lambda, and ranges over all possible times. The truth-conditional contribution of the λ -abstractor is as always:

$$[[\lambda d_0. p]]^{c, g, w} = \lambda t_0. [[p]]^{c, g(t_0/d_0), w}$$

This derives the right truth-conditions for (5b): it is true at c , g and w iff for all w_1 and t_1 compatible with what Agnes believes in w and c_T , it rains in w_1 and t_1 . Thus, Agnes’s belief needn’t be *de se*: it does not make any special reference to the present time. This also works for the past time; consider the following LF for (5c):

(5c) Agnes believes that it rained

LF of 5c: d_0^{pres} Agnes believes λd_1^{pres} . d_2^{past} [$d_2 < d_1^{pres}$] rain

Again, “rained” contributes time variables with tense features, but now the present tense feature is erased because it is bound by “believes”. This yields the desired result, as Agnes’s belief is not *de se*: her belief is not about a moment previous to c_T , but about times in general. Her belief in c and w is the set of worlds w_1 and times t_1 such that it rains at w and t_2 , under the presupposition that t_2 occurs before t_1 . This also accounts for the following contrast:

(5d) Agnes thought it rained

(5e) Agnes thought it had rained

In (5d) we do not want the content of Agnes’s thought to be about a time previous to her thinking: the past tense of the embedded sentence should play no role, unlike in (5c). In (5e) we do want Agnes’s thought to be about a time previous to her thinking. This can be easily achieved if we apply FTUB and treat “had” as an object-language existential quantifier:

LF of (5d): d_1^{past} [$d_1 < d_0^{pres}$] Agnes thinks λd_2^{past} . d_2^{past} rain

LF of (5e): d_1^{past} [$d_1 < d_0^{pres}$] Agnes thinks λd_2^{past} . [$\exists d_3$: $d_3 < d_2^{past}$] rain

Again, the past features of “rain” are erased when it is bound by a past tense verb like “rain”; this erases the presupposition and yields Agnes’s belief tenseless in (5d). Again, this yields the desired result: (5d) is true at c , g and w iff, for all worlds w_2 and times t_2 compatible with Agnes’s beliefs at $g(d_1)$ and w , it is raining at w_2 and t_2 , where $g(d_1)$ takes place before c_T . (5e), on the other hand, is a bit different: “had” can be seen as contributing an existential quantifier ranging over times and whose domain of quantification is restricted to times occurring before the time of the verb

it is attached to. Therefore it is true at c , g and w iff, for all worlds w_2 and times t_2 compatible with Agnes's beliefs at $g(d_1)$ and w , there is a time t_3 occurring before t_2 and it is raining at w_2 and t_3 , where $g(d_1)$ takes place before c_T . I.e., Agnes's belief in (5d) is represented as a set of circumstances in which it rains, whereas in (5e) it is a set of circumstances such that it rained before the time of the circumstance.

Now we have the resources at hand for dealing with the cases we are concerned with. All that is left to do now is simply to specify the following lexical entries for the relevant temporal operators. "Two days earlier" can be simply thought of as contributing to the past presupposition:

LF for "it rained two days earlier": $d_1^{\text{past}} [d_1 < d_0^{\text{pres}} \text{ and } d_1 \text{ occurs two days earlier than } d_0]$ rain

Giving to this clause its obvious presuppositional semantics, a number of facts can be explained, such as why "two days earlier" cannot be combined with future tense (the presupposition would be contradictory). It also yields the right results for (5a), i.e., for "Agnes said it had rained two days earlier":

LF for (5): $d_1^{\text{past}} [d_1 < d_0^{\text{pres}}]$ Agnes says $\lambda d_2^{\text{past}}. [\exists d_3: d_3 < d_2^{\text{past}} \text{ and } d_3 \text{ occurs two days earlier than } d_2^{\text{past}}]$ d_3 rain

I.e., the content of Agnes's saying at $g(d_1)$ and w can be thought of as the set of worlds w_1 and times t_2 such that there is a time t_3 occurring 2 days before t_2 and it rains in t_3 and w_1 . And now let us go for the crucial point: the lexical entry for "two days ago". We may treat it as contributing a new variable, but this variable is to be treated just like PRO: it is a logophor and it must be bound. Namely, this variable should be λ -bound by a previously appearing tense variable; that aside, its contribution can be just like that of "two days earlier". We may provide an LF like the following, where d_2 is the variable contributed by "two days earlier":

LF for (5a): $d_1^{\text{past}} [d_1 < d_0^{\text{pres}}]$ $\lambda d_2. d_2$ Agnes says $\lambda d_3^{\text{past}}. [\exists d_4^{\text{past}}: d_4 < d_3^{\text{past}} \text{ and } d_4 \text{ occurs two days earlier than } d_2^{\text{past}}]$ d_4 rain

This means that the past time variable d_1 , contributed by "said", works as if it were raised. Since the lambda to which it is attached quantifies into the embedded sentence, it has a role to play in the determination of the proposition reported: (5b) is true at c , g and w iff $g(d_1)$, which occurs before c_T , is a time t_2 such that, for all possible worlds w_1 and times d_3 compatible with what Agnes says at w and t_2 , there is a time t_4 taking place before t_3 and also occurring 2 days before t_2 (this is the crucial part) such that it rains in t_4 and w_1 . A bit cumbersome, but it predicts the right truth-conditions for the *de se* reading: the content of Agnes's assertion at $g(d_1)$ is the set of circumstances consisting of a possible world w_1 and a time t_4 such that t_4 takes place 2 days before $g(d_1)$ and it rains at w_1 and t_4 . That's why it is true only in situation E, in which Agnes is aware that the rain occurred 2 days before the time of her utterance.

If all these claims are right, then traditional semantics has the resources for dealing with the phenomenon of the *de se* in attitude reports: all the cases of *de se* awareness in reporting attitudes can be analyzed as simply cases of variables bound

at the level of LF. Propositions (intensions) are not fine-grained enough, but this does not mean that we should enrich them nor reject them: binders are hyperintensional operators, which means that they operate on a level that goes beyond propositions.

Now, an objection to this idea is immediately at hand: aren't there sentences in which there is binding but no clear *de se*? I shall discuss this objection in the next subsection.

5.4 A Final Objection

Chierchia (1989: 10) complains that binding alone cannot do all the job. Consider the following sentence:

(12) Everyone in this room thinks that she is in danger

Normally, (12) would be read *de se*. However, one can devise some kind of weird situation in which (12) must be read obligatorily *de re*: imagine, for example, that the room referred to in (12) is full of mirrors, and that each of the persons inside it is looking at her own reflection without being aware that she is looking at herself. Suddenly, a group of bears enters the room, and everyone inside it sees that her own reflection is about to be attacked by one of these animals. In such a situation – I am sorry for its weirdness – each of the persons in the room will have a *de re* but non-*de se* belief.

Chierchia's solution is very easy: on the *de re* reading, "think" relates each person x in the room to a proposition (the proposition that x is in danger), whereas on the *de se* reading it relates each of them to a property:

De re (Chierchia): [[[Everyone in this room] $\lambda i. t_i$ think he_i is in danger]]^g = 1 iff

$[\lambda P. \text{Every } x (x \text{ is in this room} \rightarrow P(x)) (\lambda y. y \text{ think } y \text{ is in danger})$

De se (Chierchia): [[[Everyone in this room] $\lambda i. t_i$ think $\lambda j. he_j$ is in danger]]^g = 1 iff

$[\lambda P. \text{Every } x (x \text{ is in this room} \rightarrow P(x)) (\lambda y. \text{think } \lambda z. z \text{ is in danger})$

However, this objection isn't hard for us to overcome. All we need is to replicate the binding so that, for each of the persons in the room, the content of her thought has nothing to do with herself. Namely, we have to raise the subject of the embedded sentence. My analysis of the *de re* reading is just as Chierchia's, but my analysis of the *de se* differs (again, I will be omitting relativization to contexts and circumstances):

De se (mine):

LF: [Everyone in this room] $\lambda i. t_i$ thinks that [he_i] $\lambda j. t_j$ is in danger

[[[Everyone in this room] $\lambda i. t_i$ thinks that [he_i] $\lambda j. t_j$ is in danger]]^g = 1 iff

[[Everyone in this room]]^g ([[$\lambda i. t_i$ thinks that [he_i] $\lambda j. t_j$ is in danger]]^g) = 1 iff

[[Everyone in this room]]^g ($\lambda x. [[\text{thinks}]^{\text{g}(x/i)} ([[t_i]]^{\text{g}(x/i)}, [[[\text{he}_i] \lambda j. t_j \text{ is in danger}]]^{\text{g}(x/i)})]] = 1$ iff

[[Everyone in this room]]^g ($\lambda x. \text{thinks}(x, [[[\text{he}_i] \lambda j. t_j \text{ is in danger}]]^{\text{g}(x/i)}) = 1$ iff

[[Everyone in this room]]^g ($\lambda x. \text{thinks}(x, [[[\text{he}_i]]^{\text{g}(x/i)} ([[\lambda j. t_j \text{ is in danger}]]^{\text{g}(x/i)})]] = 1$ iff

[[Everyone in this room]]^g ($\lambda x. \text{thinks}(x, [\lambda P. P(x)] ([[\lambda j. t_j \text{ is in danger}]]^{\text{g}(x/i)}) = 1$ iff

[[Everyone in this room]]^g ($\lambda x. \text{thinks}(x, [\lambda P. P(x)] (\lambda y. [[t_j \text{ is in danger}]]^{\text{g}(x/i)(y/j)}) = 1$ iff

[[Everyone in this room]]^g ($\lambda x. \text{thinks}(x, [\lambda P. P(x)] (\lambda y. y \text{ is in danger})) = 1$ iff

[[$\lambda Q. \text{Every } z (z \text{ is in this room} \rightarrow Q(z)) (\lambda x. \text{thinks}(x, [\lambda P. P(x)] (\lambda y. y \text{ is in danger})) = 1$ iff

Look at the seventh line in the calculation: the assignment is shifted for the second time. Thus, every x will have a thought whose content does not have anything to do with x herself. Or, to put it otherwise, what is believed by every x is not the content that x is in danger but, again, a function from assignments to propositions.

This discussion brings us to a final *caveat*. I have argued that *de se* readings of a sentence are just cases of variable binding, but the converse is obviously not true: not all cases of binding trigger *de se* readings. The *de re* reading of (12) is an obvious case, but some additional cases include (8) and (10), which can hardly fit within the traditional model of *de se* attitudes. Of course, this disanalogy also arises for the other theories discussed – Chierchia, for example, always had in mind that having a relation of individuals to properties is *necessary but not sufficient* to license this kind of attitudes (Chierchia 1989: 17). What this means is that some additional criteria are required for identifying which cases of binding can be considered as instances of *de se* readings of a sentence; in the case of indirect discourse, the occurrence of a propositional attitude verb seems to be a *prima facie* reasonable requirement, but probably there are many other criteria that have to be met. Be it as it may, the establishment of these criteria lies beyond the scope of this essay.

6 Conclusion: A Dilemma

Throughout this essay I have been arguing that all the inaccuracies attributed to Kaplan's LD regarding the context-sensitivity of indirect discourse are just cases of variable binding. Thus, all the previous adjustments to the semantics seem unnecessary: there is no need to enrich the circumstance with an agent, and attitudinal verbs needn't undergo type-shifting so as to be able to take either propositions or properties. All the resources that I have used are traditional and

standard in semantics. However, my final result does not seem quite Kaplanian, since, as we have seen, every *de se* attitude contains a monster. Moreover, a consequence of my view is that a *de se* report will never share its content with the original utterance: the content of the embedded sentences in (3) or (3a) is not the same as the content of Agnes's utterance of "I will be elected".

One possible way out of this is to depart from Kaplan by adding the assignment function as a part of the circumstance, so that the content of a sentence is a function from worlds, times and assignments to extensions. If we do this, then variable binders are not monsters any more, since sentences containing them will be compositional at the level of content – basically, assignment shifters would have the same status as modal and temporal operators. This idea, however, does not seem quite Kaplanian either, since, in this case, demonstratives cannot be considered as directly referential devices: their reference is mediated by the assignment. Moreover, it does not make much sense to say that what is said is true or false depending on a world, a time and an assignment function.

But I think that the results here can be made satisfactory if we see them under another light. Traditional approaches to semantics, as argued in Sect. 5.2, cannot be compositional at the intensional level, since they require lambda abstractors, which do not operate on (traditional) intensions. Frege's problem, in turn, was also a problem of compositionality, since he wanted to derive the meaning of "Gottlob believes that Hesperus is Hesperus" and "Gottlob believes that Hesperus is Phosphorus" compositionally. This may be symptomatic: perhaps compositionality is not to be sought at the intensional level. After all, monsters are compatible with compositionality at the hyperintensional level (see Westerståhl 2012 and Briciu 2018 for proofs), and no one vindicates compositionality at the extensional level nowadays, so a further step in this direction should not be so harmful.

Acknowledgements I wish to thank Ramón García Moya, Stefano Predelli, Jordi Valor and audiences in Barcelona and Warsaw for their comments on previous versions of this paper.

References

- Anand, P., & Nevins, A. (2004). Shifty operators in changing contexts. *Proceedings of SALT, 12*, 20–37.
- Briciu, A. (2018). On context-shifters and compositionality in natural languages. *Organon, 25*(1), 2–20.
- Cappelen, H., & Dever, J. (2013). *The inessential indexical: On the philosophical insignificance of perspective and the first person*. Oxford: Oxford University Press.
- Castañeda, H. (1966). 'He': A study in the logic of self-consciousness. *Ratio, VIII*(2), 130–157.
- Castañeda, H. (1967). Indicators and quasi-indicators. *American Philosophical Quarterly, 4*(2), 85–100.
- Chierchia, G. (1989). Anaphora and attitudes *de se*. In R. Bartsch, J. van Benthem, & B. van Emde (Eds.), *Language in context* (pp. 1–31). Dordrecht: Foris.
- Clements, G. (1975). The Logophoric pronoun in ewe: Its role in discourse. *Journal of West African studies, 10*(2), 141–177.

- Crossley, J., & Humberstone, L. (1977). The logic of “actually”. *Reports on Mathematical Logic*, 8, 11–29.
- de Ponte, M., Korta, K., & Perry, J. (2019). Utterance and context. In T. Ciercierski & P. Grabarczyk (Eds.), *The architecture of context and context-sensitivity*. Berlin Heidelberg: Springer.
- Heim, I. (2001). *Features of pronouns in semantics and morphology* (Manuscript).
- Heim, I., & Fintel, K. (2011). *Intensional semantics* (Manuscript).
- Heim, I., & Kratzer, A. (1998). *Semantics in generative grammar*. Oxford: Blackwell.
- Kaplan, D. (1989a). Demonstratives. In J. Almog et al. (Eds.), *Themes from Kaplan* (pp. 481–564). Oxford: Oxford University Press.
- Kaplan, D. (1989b). Afterthoughts. In J. Almog et al. (Eds.), *Themes from Kaplan* (pp. 565–614). Oxford: Oxford University Press.
- Lewis, D. (1979). Attitudes *de dicto* and *de se*. *The Philosophical Review*, 88(4), 513–543.
- McCullagh, M. (2017). Kinds of monsters and kinds of compositionality. *Analysis*, 77(5), 1–11.
- Montague, R. (1973). The proper treatment of quantification in ordinary English. In J. Hintikka, J. Moravcsik, & P. Suppes (Eds.), *Approaches to natural language* (pp. 221–242). Dordrecht: Reidel.
- Ninan, D. (2010). *De Se* attitudes: Ascription and communication. *Philosophy Compass*, 5(7), 551–567.
- Ninan, D. (2016). What is the problem of *de se* attitudes? In M. García-Carpintero & S. Torre (Eds.), *About oneself: de se attitudes and communication* (pp. 86–120). Oxford: Oxford University Press.
- Perry, J. (1977). Frege on demonstratives. *The Philosophical Review*, 86(4), 474–497.
- Perry, J. (1979). The problem of the essential indexical. *Nous*, 13(1), 3–21.
- Predelli, S. (2014). Kaplan’s three monsters. *Analysis*, 74, 389–393.
- Prior, A. (1968). Now. *Nous*, 2, 101–119.
- Rabern, B. (2012). Against the identification of assertoric content with compositional value. *Synthese*, 189, 75–96.
- Rabern, B. (2013). Monsters in Kaplan’s LD. *Philosophical Studies*, 164, 393–404.
- Rabern, B. (2014). *Being a monster* (Manuscript).
- Rabern, B., & Ball, D. (2019). Monsters and the theoretical role of context. *Philosophy and Phenomenological Research*, XC VIII(2), 392–416.
- Schlenker, P. (2002). Indexicality, logophoricity, and plural pronouns. In J. Lecarme (Ed.), *Afroasiatic grammar II: Selected papers from the fifth conference on afroasiatic languages, Paris, 2000* (pp. 409–428). Amsterdam: John Benjamins.
- Schlenker, P. (2003). A Plea for monsters. *Linguistics and Philosophy*, 26(1), 29–120.
- Stalnaker, R. (1981). Indexical belief. *Synthese*, 49, 129–151.
- von Stechow, A. (2002). Feature deletion under semantic binding: Tense, person, and mood under verbal quantifiers. *Proceedings of NELS*, 33, 379–403.
- Westerståhl, D. (2012). Compositionality in Kaplan-style semantics. In M. Werning, W. Hinzen, & E. Machery (Eds.), *The Oxford handbook of compositionality*. Oxford: Oxford University Press.

How Can “I” Refer to Me? Banishing Monsters at the Source



David Kashtan

Abstract Kaplan’s influential (1989) makes a connection between the mode of reference of indexical expressions and the impossibility of certain sentential operators, which he calls monsters. The impossibility of monsters has recently come under attack from several quarters, both theoretical and empirical. In this paper I consider monsters from a different perspective. I motivate the prohibition on monsters independently of intensional notions, and understand it not as an empirical hypothesis, but as an adequacy criterion on formal systems intended to capture the philosophical notion of direct reference. I show that Kaplan’s formalism doesn’t live up to this criterion, and sketch in preliminary fashion a formalism that does.

Keywords Semantics · Formal semantics · Metasemantics · Pragmatics · Semiotics · Reference · Indexicals · Direct reference · Semantic monsters

Kaplan’s influential (1989) makes a connection between the mode of reference of indexical expressions and the impossibility of certain sentential operators, which he calls *monsters*. The impossibility of monsters has recently come under attack from several quarters. Rabern (2013), for example, defines monsters in terms of intensional compositionality, and argues that they are inevitable already in extensional semantics.¹ Schlenker (2002), and others following him, bring forth empirical evidence for the existence of monsters in natural language.² On the

¹See also Rabern and Ball (2019). Linking the prohibition on monsters to content compositionality is common, see Westerståhl (2012), Stalnaker (2014), McCullagh (2017), Briciu (2018).

²See also Schlenker (2011), Anand and Nevins (2004). But see Stalnaker (2014, 210) for an objection.

D. Kashtan (✉)

Edelstein Center for the History and Philosophy of Science, Hebrew University of Jerusalem, Jerusalem, Israel

defending side, Predelli (2014) disentangles three senses of “monster”, hoping that some senses can be sacrificed to save the others.

In this paper I wish to consider monsters in a slightly different perspective. Contra Rabern, I wish to motivate the prohibition on monsters independently of the requirement for content compositionality, indeed independently of intensional notions altogether. And contra Schlenker, I wish to understand the prohibition on monsters not as an empirical hypothesis, but as an adequacy criterion on formal systems intended to capture the philosophical or semiotic notion of direct reference. These two tasks will be discharged in §1. In §2 I’ll rehearse the formal structure of Kaplan’s system, and show that monsters are definable on it. This means that Kaplan’s own formal system fails to capture his philosophical thesis about direct reference. §3 motivates and develops, in rough outline, an alternative formalism designed to capture the special behavior of indexicals without giving rise to monsters.³

1 Direct Reference and the Prohibition on Monsters

Indexicals are usually characterized as a class of context-dependent natural language expressions. “Context”, however, is a theory-laden term, so for the moment we’ll think of indexicals intuitively as the class of the singular terms “I”, “now”, “you”, etc.⁴ Kaplan’s philosophical thesis about indexicals consists of the following two principles (1989, 492):

Context-dependence. *Indexical reference is context-dependent.*

Direct reference. *Indexical singular terms refer to their objects directly.*

These principles have a corollary that Kaplan introduces with the following words (510, emphasis in the original):

Are there such operators as ‘in some contexts it is true that’, which when prefixed to a sentence yields [sic] a truth if and only if in some context the contained *sentence* . . . expresses a content that is true in the circumstances of that context? Let us try it:

³This work stems from my unpublished doctoral dissertation, Kashtan (unpublished), and departs from it somewhat. I would like to thank the following for their role in helping the work reported in this paper to come to existence, whether through support, helpful criticism or otherwise: Moysh Bar-Lev Eli Dresner, Rea Golan, Michael Goldboim, Balthasar Grabmayr, Aviv Keren, Ran Lanzet, Carl Posy, Danny November, Gil Sagi; an audience at the Language, Logic and Cognition Center seminar at HUJI (esp. Edit Doron), an audience at the Context, Cognition and Communication 2018 conference in Warsaw; and especially Paweł Grabarczyk and Tadeusz Ciecierski for organizing the CCC conference and getting this volume together.

⁴Kaplan’s “pure indexicals”.

[*] In some contexts it is true that I am not tired now.

For [*] to be true in the present context it suffices that some agent of some context not be tired at the time of that context. [*], so interpreted, has nothing to do with me or the present moment. But this violates [the direct reference thesis]! . . . I am not saying we could not construct a language with such operators, just that English is not one. And such operators *could not be added to it*.

Following Kaplan we will say, loosely at this stage, that an operator (e.g. an adverb) that affects the reference of the indexical in its scope is a *monster*. Understood simply, the quote above can be summarized as here:

The prohibition on monsters. *There are no monsters in English.*

The two principles and the corollary constitute, as far as this paper is concerned, the core of Kaplan’s philosophical theory of indexicals. Three questions present themselves: What exactly does “direct reference” mean? How is direct reference related to context dependence? Finally, how is the prohibition on monsters derived from the two principles?

1.1 From Direct Reference to the Prohibition on Monsters

We can answer these questions if we find some underlying notion that ties the principles and corollary together, yielding a kind of “metasemantic” characterization of direct reference.⁵

1.1.1 Grounding Direct Reference

What does it mean for an expression to refer *directly*? In Kaplan, this idea is explicated mainly in negative terms, through its contrast with reference *mediated* by an intensional or descriptive layer of meaning, called by Kaplan *content*. The “direct” vs. “mediated” idiom is metaphorical, however, and it is more to the point to ask, for each mode of reference, what it is *grounded* in. In the case of content-mediated reference we have a short answer. Consider, for example, the non-indexical singular term “the first US president”. This term refers to the unique individual which satisfies the general term “first US president”. It is therefore in virtue of certain satisfaction conditions being fulfilled that the term refers to what it refers. Descriptive reference is grounded in the concepts of satisfaction and truth.

We want to tell a similar story for indexical reference. Kaplan’s own metaphysics of direct reference relies heavily on a picture of structured propositions and possible worlds. For example, Kaplan describes direct reference as one which contributes

⁵Roughly in Burgess and Sherman’s (2014, 2) sense, of metasemantics as “the business of providing metaphysical explanation of semantic facts”.

the referent itself to the proposition it is part of, rather than a complex of attributes (p. 494). Alternatively, Kaplan speaks of “obstinate rigid designation” (p. 493), which in turn is explicated in terms of possible worlds.⁶ But these are notions from the domain of descriptive reference.⁷ It would be better to explicate direct reference in independent terms.

We look to etymology for a clue. The term “indexical”, as Kaplan uses it, comes from Bar-Hillel (1954), who places the study of indexicals under the rubric of pragmatics.⁸ Bar-Hillel has in mind Carnap’s (1942, 8f) division of the study of language, or semiotics, into syntax, semantics and pragmatics, where syntax studies relations between expressions, semantics relations between expressions and their referents, and pragmatics relations between expressions, referents and users.⁹ This suggests that indexical reference is grounded by some relation between expression, denotation and user. But which relation?

Bar-Hillel borrows “indexical” from C.S. Peirce’s theory of signification, or semiotics.¹⁰ Peirce classifies signs into three groups: icons, symbols and indices, based on their mechanism or mode of signification. Icons signify in virtue of a similarity between signifier and signified; symbols in virtue of arbitrary convention; and indices in virtue of what Peirce calls “existential connection”, sometimes “physical contiguity”, between signifier and signified. Unlike the other two modes of signification, indexical signification depends on the physical existence of the sign and on its spatio-temporal situation, or in Kaplan’s terms, its “utterance”.¹¹ It is the physical, situated, event of the utterance of an indexical that does the referring, unlike a descriptive phrase such as “the first US president”, which refers to the same thing regardless of when and whether it is uttered. Consequently, the referent of an indexical utterance will be an object which stands in some “existential connection” to the utterance, by which is meant some kind of physical relation. Assuming that “existential connection” is transitive, and that an utterance always stands in existential connection to the utterer, we have our first candidate for a referent-user relation grounding indexical reference.

An alternative story comes from a more proximal source of inspiration for Kaplan, Russell’s distinction between reference by description and reference by acquaintance.¹² According to Russell, proper names can refer, in the strict sense of the word, only to things with which their user has direct acquaintance; in the

⁶See Roberts (1994) for a discussion of the tension between Kaplan’s two metaphysical pictures.

⁷Kaplan (1989, 493fn17) remarks on that.

⁸See also Montague (1970).

⁹Carnap takes up this division from Charles Morris.

¹⁰This is the source of Morris’s term “semiotics”, but used in a different meaning. See Atkin (2005) for Peirce’s view of indexicals.

¹¹I’ll be abstracting from various issues such as non-conversational uses of language (see Cohen and Michaelson (2013) for an overview) or the token vs. utterance vs. sentence-in-context debate (see García-Carpintero (1998) for an example).

¹²See Russell (1905).

absence of direct acquaintance, a proper name serves only to assert the existence of a (unique) object with the described properties, and this falls short of reference proper.¹³ On this view, the relation between the user and the referent of an indexical, in virtue of which indexicals are said to refer, has an irreducible cognitive aspect (acquaintance) which is missing on the Peirce-inspired account.

I won’t try to adjudicate between these roughly sketched alternatives. It will suffice to assume that there is some relation between user and referent that grounds indexical reference. I’ll call this relation *presence* and treat it as a primitive term. This isn’t a simple binary relation, since typically many things will be present, but only one will be referred to though a given indexical. For example, though both you and I are present to me now (let’s assume), the indexical “you” refers to you and not to me. Hence presence is somehow modulated and this is something indexical reference is sensitive to. We can think of these “modes of presence” as positions relative to the speaker. I’ll represent relative positions as mere labels (strings), i.e. “speaker”, “addressee”, etc.

The first of our three guiding questions from the top of the section was about the nature of direct reference. Our answer is that direct reference is reference grounded in (modulated) presence, a notion not belonging to the domain of descriptive reference. The direct reference thesis boils down to the claim that (the utterance of) an indexical may only refer to a present object.

1.1.2 Deriving the Prohibition on Monsters

Taking presence as our primitive term, we can define the *context* of an utterance to be the labeled collection of objects present to the agent.¹⁴ The modes of presence are then called *context features*. Kaplan’s context-dependence principle now becomes a definitional variant of the direct reference thesis. This answers the second of our three questions, about the relation between the two principles.

The third question was about the derivation of the prohibition on monsters from the direct reference thesis. To show it, we need to make explicit an assumption that Kaplan is committed to implicitly. Observe the sentence:

(1) The president lives in the White House.

We say that the phrase “the president” in (1) refers to a specific individual.¹⁵ By contrast, in (2) there is no specific individual to which “the president” refers.

(2) It is always the case that the president lives in the White House,

However, the evaluation of (2) is often understood in terms of a multiple evaluation of its subclause (1), one case of evaluation for each instant covered by

¹³Russell says it “denotes”. I won’t be fussy about terminology.

¹⁴I leave “collection” vague.

¹⁵*Pace* Russell, see footnote 10.

“always”. In each such evaluation of (1), the reference of “the president” is evaluated relative to corresponding instant. Relative to 2016, for example, “the president” in (1) refers to Obama; relative to 2018 it refers to Trump; and so on. Thus although “the president” in (2) doesn’t strictly speaking refer to any one object, we can say that *in a secondary sense* it refers both to Obama and to Trump, and to every past, present or future president of the US.

Now assume, as Kaplan implicitly does, that the direct reference principle applies to reference also in this secondary sense. Then we can derive from it the prohibition on monsters, as follows. Suppose that yesterday John uttered truly the sentence:

(3) I am tired.

Now assume that there is a monstrous operator “*M*” such that a sentence such as

(4) *M* (I am tired),

is true just in case some speaker in some context is tired. Then an utterance by me of (4) will be true in virtue of the truth of John’s utterance of (3) yesterday; in particular the indexical “I” in my utterance of (4) will refer (in the secondary sense) to John, even if John is not present in the context. But this violates The direct reference thesis. More generally, any linguistic expression the intended effect of which is to make the utterance of an indexical refer (in the secondary sense) to something absent from the context of utterance is ruled out by Kaplan’s theory of direct reference. And this is just the prohibition on monsters.¹⁶

As formulated above, the prohibition on monsters is an observable consequence of the metaphysical thesis of direct reference. As such, it can be used in an empirical test for the latter:

Kaplan’s Descriptive Criterion. *If monsters exist then the direct reference thesis is false.*

1.2 From the Prohibition on Monsters to Direct Reference

1.2.1 Schlenker’s Challenge

The prohibition on monsters has been challenged in the past on theoretical grounds.¹⁷ More recently, empirical evidence has been adduced against it. Schlenker (2002), and others after him, claim to have found clear cases of monstrous operators

¹⁶Compare this with, e.g., Rabern (2013), who ties the prohibition on monsters to content-compositionality. See also Rabern and Ball (2019), who explain it in terms of structured propositions. On my account, direct reference and the prohibition on monsters are understood independently of any prior conception or claim about content and intensionality.

¹⁷See Israel and Perry (1996).

in natural language.¹⁸ For example, Schlenker locates a sentence in Amharic that, when glossed word for word, reads:

(5) John says that I am a hero,

but such that the word glossed by the indexical “I” refers to John, even in contexts in which John is not present. Schlenker’s conclusion is that “say”, and other attitude verbs, are monsters. Assuming that this evidence is sound,¹⁹ by the criterion of the previous section, Kaplan’s thesis of direct reference is refuted, falsified by the empirical evidence.

But is Kaplan’s prohibition on monsters really meant as an empirical prediction? Such a reading seems to be justified by the first part of Kaplan’s statement (p. 510):

I am not saying we could not construct a language with such operators, just that English is not one.

Reading “natural language” or “Amharic” for “English”, this statement is refuted by (5). However, Kaplan then goes on to say (emphasis in the original):

And such operators *could not be added to it*.

The modal in this sentence lends it an aprioristic tone which is hard to reconcile with the empirical-sounding claim of the previous sentence. Natural language, presumably, is an empirical phenomenon, the properties of which are for us to discover, not to decide. What could it mean to *add* an operator to a natural language, and what does it mean for an operator to be *impossible* to add?

1.2.2 Two Semiotic Projects

In Kaplan (1989) I don’t find anything that helps me to resolve this ambivalence. If you’ll allow me a brief historical aside, this is not accidental. Kaplan was writing in a transitional period, in which formalized languages were coming to be regarded as devices for modeling natural language phenomena, in contrast to the more traditional outlook of philosophers such as Frege, Tarski and Quine, who thought of them rather as a recommended medium for scientific discourse. The two approaches are disparate enough in their goals and methods to be called two separate projects. Carnap (1942, 11f) is quite clear about this distinction (emphasis in the original):

Descriptive semantics is the empirical investigation of the semantical features of historically given languages. [*Pure semantics* is] the construction and analysis of system[s] of semantic rules, [which] constitute nothing else than a definition of semantical concepts such as designation and truth. It is entirely analytic and without factual content. We make an analogous distinction between *descriptive* and *pure syntax*.

¹⁸See Schlenker (2011) for an overview with later references.

¹⁹See von Stechow (2003), Stalnaker (2014, 211) for arguments to the contrary.

Following Carnap, I'll make a distinction between a *descriptive* and a *foundational* project in semiotics or the philosophy of language.²⁰ Descriptive semiotics starts with a body of data and proceeds to formulate a theory that explains it, in the sense of predicting it, or, in a terminology more proper to linguistics, generating it. This data varies with the precise project being engaged in: it can be a set of well-formed sentences if syntax is our target, or truth conditions and entailment relations if we're interested in semantics, etc. The point of the foundational project, by contrast, is to make precise sense of certain fundamental semiotic notions. Carnap's paradigm for pure semantics is Tarski's work on the notions of truth and of logical consequence. Tarski isn't interested in faithfully describing facts of usage of the adjective "true" and the construction "logically entails".²¹ He is interested in clarifying the philosophical concepts that go under these labels, independently of whether his clarification matches up with ordinary usage or natural language.²²

The difficult question about the foundational project is to say how we know whether a proposed analysis is successful. Descriptive adequacy is relatively straightforwardly explained in terms of adequacy to the data, but on the foundational project there is nothing analogous to data. This problem deserves more discussion than will be offered here. Taking a cue from Tarski on truth, we'll take agreed-upon formal criteria of adequacy, in the likeness of Convention T, to be the basis and the starting point of any foundational investigation. Once an adequacy criterion is established, the job is to come up with an analysis of the investigated notion that is formulated in terms precise enough to allow the application of the criterion.

Indexicality, according to Bar-Hillel, is part of pragmatics. Carnap himself doesn't make a pure/descriptive division for pragmatics. In fact, several pages forward he says that "descriptive semantics and syntax are, strictly speaking, parts of pragmatics".²³ The reason is that in order to study semantics and syntax descriptively, we must study the behavior of users, which is a pragmatic notion. This suggests that pure pragmatics is impossible. But this is a mistake. Granted, pragmatics is about users, and descriptive syntax and semantics involves the study of use. But the notion of use, like those of reference and truth, can be approached either descriptively or foundationally. Indeed, Bar-Hillel (1954, 369) recommends "the erection of indexical language-systems" as a task of "pure pragmatics".

²⁰I prefer "foundational" because "pure" has unwanted axiological connotations.

²¹*Pace*, e.g., Gupta (2012, 54).

²²See Stalnaker (2014) for a similar distinction. Stanley and Szabo (2000, 220ff) make a distinction with the same titles, but different content. One sometimes distinguishes between descriptive and theoretical linguistics. Both sides of this distinction fall on the descriptive side of Carnap's distinction, and correspond, respectively, to Carnap's (1942, 11) *special* vs. *general* semiotics.

²³(1942, 13). See also (1955).

1.2.3 Kaplan’s Criterion Again

We saw that Kaplan, in the monsters passage, vacillates between empirical and aprioristic idiom. I suggest that this is due to a blurring in Kaplan’s mind (and in Kaplan’s generation) of Carnap’s distinction. Redrawing it, we disentangle two possible claims in Kaplan. The first, labeled *Kaplan’s descriptive criterion* above, is an empirical prediction about natural language, which is falsified by Schlenker’s evidence. The other way to understand Kaplan’s text is as proposing a criterion of adequacy on formal semantic systems that purport to capture the metaphysical notion of direct reference. On this understanding, the prohibition on monsters isn’t, and can’t be, falsified by any kind of empirical evidence, though it might turn out that natural languages do not contain any indexicals.²⁴ On this approach, Kaplan’s theory is a chapter of foundational pragmatics.

Kaplan’s Foundational Criterion. *If monsters are definable in a formalism, then that formalism does not capture direct reference.*

If you agree to accept this construal of Kaplan’s criterion, then the direct reference thesis is rescued. However, we will now see that Kaplan’s own formal system fails to conform to it.

2 Précis of Formal Semantics

Alongside the metaphysics of direct reference, Kaplan offers a formal system designed to capture it, which has become the standard formalism for indexical semantics. However, by the above criterion, it is foundationally inadequate, a fact concealed by the general blurring of the descriptive/foundational distinction. In this section I will analyze Kaplan’s semantics (a simplified version thereof) into its three aspects: extensional, intensional and indexical, and look into each aspect from the foundational perspective.

2.1 Extensional Semantics

I will restrict attention to first-order languages with the usual logical vocabulary. Such languages differ in their non-logical vocabulary (their “ideology”, in Quinean) and in their domain of quantification (“ontology”). As an example, let L_x be a standard first-order language whose non-logical vocabulary consists of a unary predicate Px (holding of presidents) and an individual constant a (referring to

²⁴Let me note that Schlenker (2002, 29) is not blind to this option, but doesn’t find it attractive. If you’re skeptical about the possible intrinsic interest of expressive devices that can’t be found in natural language, recall Kaplan’s (1978) own *dthat*.

Trump). For concreteness, let the domain of quantification D^{L_x} include all people, living or otherwise. Typical L_x sentences include Pa , $\forall x(Px \rightarrow x = a)$, etc.²⁵

The semantic theory for a language like L_x can take the form of a denotation function, $\llbracket x \rrbracket$. For example:

$$(6) \quad \llbracket a \rrbracket = Trump,$$

$$(7) \quad \llbracket P \rrbracket = \lambda x. (1 \text{ iff } x \text{ is president}) \cong \{x : x \text{ is president}\},$$

$$(8) \quad \llbracket Pa \rrbracket = 1 \text{ iff } Trump \text{ is president.}$$

The double-bracket notation is familiar and convenient when the denotation function is fixed, but it can be confusing when comparing several denotation functions, as we will do. Consequently, I will rather write $xden(e)$ for the denotation of the expression e , understood as an expression of an extensional first-order language.

The denotation function $xden(e)$ of a language L_x may be defined in a suitable metalanguage M_x , in ways based on Tarski (1933). The definition is a single sentence which mentions the domain D^{L_x} and the lexicon. Let $base_{L_x}$ be a metalinguistic sentence specifying the lexicon, e.g.

$$(9) \quad base_{L_x} = "xden(a) = Trump \wedge xden(P) = \{x : x \text{ is president}\}."$$

Definability. Let " $def_x(base_{L_x}, D^{L_x})$ " abbreviate the definition of $xden(e)$ for L_x in M_x .

Foundationally speaking, extensional languages are meant to capture the idea of a self-standing objective reality. In terms of classical grammar, their sentences are understood to be in the indicative mood, tenseless and non-indexical. Pragmatically, they are used with assertive force, i.e. they are used to make assertions about the world (broadly construed). Logically, the idea of a "self-standing" reality is captured by the defining feature of extensional semantics:

Extensional Substitutivity. $xden(s) = xden(t) \Rightarrow xden(p) = xden(p[st])$.

Here s , t and p are well-formed expressions; s and t belong to the same syntactic category; and $x[y/z]$ is the result of replacing y for z in x . This principle expresses the fact that what holds of an object or state of affairs holds of it under whatever name we give it, i.e. independently of how we represent it. The fundamental notion for extensional languages is the concept of truth, the Tarski-style definition of which yields the substitutivity principle. In the case of first-order languages, an additional notion is that of a (self-standing) object which stands in relations of satisfaction with predicate. The fundamental notion here is that of satisfaction.

Although different languages may range over radically different kinds of objects, from presidents and wars to mathematical functions and impossible states of affairs, this makes no difference from the foundational perspective. The logic is blind to

²⁵I will usually omit quote-marks.

the kind of object, being concerned only with the idea of a (self-standing) object in general. Not everything (in a loose sense of “everything”) is an object in this sense. Some things which can be talked about fail to satisfy the conditions on objecthood. The best-known example is that of proper classes in formalized set theories, for which we have a proof that they cannot coherently be included in the first-order domain.

Object-blindness. *The kind of object quantified over is foundationally irrelevant.*

2.2 *Intensional Semantics*

The development of intensional semantics is inspired by a range of phenomena from ordinary discourse which seem to exhibit failure of extensionality. Take the following pair of (almost) unregimented sentences together with their truth value (at the time of writing):

- (10) It is always the case that Trump is president (False),
 (11) It is always the case that the president is president (True).²⁶

Here the sentences “Trump is president” and “the president is president” are both true, but their embedding sentences differ in truth-value, in violation of the principle of extensional substitutivity. The failure of substitutivity is traced to the temporal adverb “always”, and is exhibited also by other kinds of expressions, most notably modal operators and propositional attitude verbs. Intensional semantics accounts for this by relativizing the denotation function to a parameter, called *index*, that represents a *circumstance of evaluation*, typically a possible world or a time.

Let L_i be a first-order language like L_x , with the addition of a sentential operator A (for “always”) and (for expository purposes) an individual constant ιP (meaning “the president”). We assume a domain of indices (times) D^I . Writing $iden(e, i)$ for the intensional denotation function, we have:

- (12) $iden(a, i) = Trump (at i)$,
 (13) $iden(P, i) \cong \{x : x \text{ is president at } i\}$,
 (14) $iden(\iota P, i) = the \text{ president at } i$,
 (15) $iden(A(\phi), i) = 1 \text{ iff } \forall i' (iden(\phi, i') = 1)$.

On intensional semantics, simple sentences are not true or false *simpliciter*, but only relative to an index. We might say that they are tensed (if the indices are times), or in the subjunctive mood (if they are possible worlds). Notice that definitions (12)–(15) have an implicit initial universal quantifier over i , which I will refer to as the *external quantifier*. An operator such as A , which is defined in terms of a

²⁶The term “the president” is meant in its “attributive use”, see Donnellan (1966).

quantifier over indices, is called an *index-shifting* (or simply *intensional*) operator. This quantifier is an *internal quantifier*. Given the semantic and the worldly facts, the following holds in L_i , capturing the contrast between (10) and (11):

$$(16) \quad \text{iden}(A(Pa), i) = 0 \neq 1 = \text{iden}(A(PtP), i).$$

In a sense, then, we have a semantics for languages without extensional substitutivity.

In addition to accounting for non-extensional operators, intensional semantics lets us capture the notion of the *content* of an expression. We do this by de-relativizing $\text{iden}(e, i)$ from the index parameter, to get a unary function from expressions to their contents, construed as mappings from indices to extensions:

$$(17) \quad \text{iden}(e) := \lambda i. \text{iden}(e, i).$$

Contents of sentences, for example, will amount to sets of indices called *propositions*,²⁷ e.g.:

$$(18) \quad \text{iden}(Pa) = (\lambda i. 1 \text{ iff } \textit{Trump is president at } i) \cong \{i : \textit{Trump is president at } i\}.$$

The intensional framework can be enriched by adding different index types and relations on the index domains. Such enriched systems can be used to account for a wide range of linguistic phenomena, including propositional attitudes, various conditionals, temporal relations, the intricacies of tense and aspect, and more.²⁸ From a foundational standpoint, however, intensional semantics is just a notational variant on extensional semantics. To see this, consider an example of a recently discussed (§1.1) modal sentence, Kaplan's statement to the effect that

monstrous operators *could not* be added to English.

According to intensional semantics as sketched here, what this statement means is that

in no possible world are monstrous operators *in fact* added to English.

This is completely unenlightening so long as we have not explicated the adjective “possible” in “possible world”. But intensional semantics does nothing of the sort. On intensional semantics “possible” is mere informal gloss. It might accurately predict the distributional properties of modal locutions in natural language, but it in no way captures the philosophical notion of possibility. Similar arguments can be made for the notions of temporality and content.

Compare this with the extensional case. What extensional languages were out to capture (according to the discussion in the previous subsection) was the idea of a reality which exists independently of how we represent it. This presumption is fulfilled (perhaps not fully) by the principle of extensional substitutivity. This

²⁷Whether times should be included in propositions or not is a question I won't discuss.

²⁸See Heim and von Stechow (2011) to get an idea.

principle is a structural, logical, property of the system, and not an informal gloss. If we construe extensional substitutivity as an adequacy condition on extensional semantics, then since Tarski’s definitions entail it, we can declare them foundationally adequate. Nothing similar, as far as I can tell, has been provided for the intensional notions. However, since my topic today is not intensionality but indexicality, I will not pursue the foundational inadequacy of intensional semantics any further.

The difference between intensional and extensional semantics comes down to the relativization of the semantics to a domain of indices, where the indices are just another kind of object like presidents and numbers. In view of the object-blindness of extensional semantics (see above), intensional semantics is better seen as just a notational variant of the latter. This is what makes it so tractable and so attractive for descriptive purposes. With straightforward modifications, we can turn the definition of an extensional denotation function $xden(e)$ to a definition of the intensional function $iden(e, i)$. The domain of indices D^I is an additional parameter:

Definability. Let “ $def_i(\text{base}_{L_i}, D^{L_i}, D^I)$ ” abbreviate the definition of $iden(e, i)$ for L_i in M_i .

As mentioned above, $iden(e, i)$ does not exhibit extensional substitutivity:

$$(19) \quad iden(s, i) = iden(t, i) \not\Rightarrow iden(p, i) = iden(p[s/t], i).$$

However, it can be established that the unary $iden(e)$ does exhibit substitutivity with regard to contents:

Intensional substitutivity. $iden(s) = iden(t) \Rightarrow iden(p) = iden(p[s/t])$.

This is shown straightforwardly from the definition of $iden(e, i)$ mentioned above. An intensional denotation function with this feature will be called *content-compositional*.

2.3 Indexical Semantics

The insight that the semantic value of indexicals should be relativized to a context parameter is recognized clearly in Bar-Hillel (1954) and implemented formally in Montague (1970). Montague used a single parameter to account both for intensionality and for indexicality. Kaplan’s central technical innovation was to distinguish between the intensional parameter, or *index*, and the indexical parameter, the *context*. Philosophically, this distinction tracks the difference between a circumstance against which an expression is evaluated (for truth, reference, etc.) and a circumstance in which an expression is used.

We stipulate a domain D^C of contexts of use and introduce a ternary denotation function $cden(e, i, c)$, which takes an expression, an index, and a context, and returns the denotation of e relative to i , as uttered in c . Metaphysically (see 1.1.2 above) we think of a context as a collection of features. Mathematically we can model this by introducing special functions, e.g. $speaker(c)$ or $time(c)$, from contexts to their

speakers or their times, respectively. We use these feature functions to define the semantic values of the indexicals:

$$(20) \quad cden(I, i, c) = speaker(c),$$

$$(21) \quad cden(now, i, c) = time(c), \text{ etc.}$$

Let L_c be the result of adding the singular terms “ I ” and “ now ” to L_i . The intensional-indexical denotation function $cden(e, i, c)$ is defined in the usual way in an appropriate metalanguage M_c , to yield truth-conditions for all sentences of the object-language, e.g.:

$$(22) \quad cden(PI, i, c) = 1 \text{ iff } speaker(c) \text{ is president at } i.$$

This is the basic idea of Kaplanian indexical semantics.

We easily get some further notions. By abstracting the index parameter i , as we did in the case of $iden(e, i)$ (§2.2.1), we can get a binary function $cden(e, c)$ that returns the content that e expresses if uttered in c . This is a kind of “content generation” function:

$$(23) \quad cden(e, c) = \lambda i. cden(e, c, i).$$

This captures the observed fact that the same sentence can express a different proposition in different contexts:

$$(24) \quad cden(PI, c) \cong \{i : speaker(c) \text{ is president at } i\}.$$

Let c_1 be a context in which Trump is the speaker and c_2 in which Hillary Clinton is. Then in c_1 , but not in c_2 , the sentence PI will express the same proposition as the sentence Pa .

If we further abstract the context parameter, we get a unary function $cden(e)$. This function returns, for each expression, its “standing” or context-independent meaning, or *character*. The character of an expression is a function from contexts to the contents the expression has in context:

$$(25) \quad cden(e) = \lambda c \lambda i. cden(e, c, i).$$

For example:

$$(26) \quad cden(I) = \lambda c \lambda i. speaker(c),$$

$$(27) \quad cden(PI) = \lambda c \lambda i. 1 \text{ iff } speaker(c) \text{ is president at } i, \text{ etc.}$$

In this way Kaplan’s simple technical innovation allows us to model a range of important notions.

By enriching the index and context domains in various ways, many linguistic phenomena can be accounted for that are not on the face of it intensional or indexical. Take imperatives for example. Various combinations of an appropriate modality (e.g. deontic), a second-person indexical, and a rich notion of context (containing, e.g., to-do-lists, preferences, etc.) have been used in giving a semantics

for imperatives. The proposals may differ significantly in their content and merit, but most of them can be brought under the basic framework of $cden(e, i, c)$.²⁹ In addition, indexical-intensional accounts can be given for presupposition, expressives, and various other phenomena.³⁰

However, on this system contexts, like indices and presidents, are just another kind of object. From a foundational perspective it should be clear that indexical semantics, like its intensional subsystem, is no more than a notational variant of extensional semantics, and for the same reasons. Indeed, we can see immediately that $cden(e, i, c)$ allows the definition of monstrous operators, and therefore that it fails Kaplan’s adequacy criterion (§1.2.3). Kaplan’s paradigmatic monster was an operator M which would make the sentence $M(PI)$ true if some speaker in some context is president. This operator can be defined in complete analogy with the temporal operator “always” (see (15) above):

$$(28) \quad cden(M(\phi), i, c) = 1 \text{ iff } \exists c' (cden(\phi, i, c') = 1).$$

It would let us utter sentences that depend for their truth on indexicals referring to an object absent from the context. For consider a context c_2 in which Hillary Clinton is the speaker. If c_1 (in which Trump is both speaker and president) is in D^C , we get:

$$(29) \quad cden(M(PI), time(c_2), c_2) = 1.³¹$$

Although the speaker of c_2 is not president, and even if nothing in c_2 is president, sentence (29) is true. This is in violation of the prohibition on monsters. Kaplan’s formal system is therefore inadequate to his philosophical thesis of direct reference.

In the next section I’ll sketch an analysis of the failure and the rudiments of an approach with more chance of succeeding. Let me conclude this section with a word about content compositionality. Some writers identify the prohibition on monsters with a requirement for content compositionality. As I argued in §1.1, this identification is wrong from a metasemantic point of view. But note that on the present framework the two conditions converge. The existence of monsters leads to violations of content compositionality. For example, the sentences $I = a$ (I am Trump) and $a = a$ (Trump is Trump) express the same proposition in any context in which Trump is the speaker:

$$(30) \quad cden(I = a, c_1) = cden(a = a, c_1) \cong \{i : \text{Trump is Trump in } i\}.$$

Let N be the monstrous operator dual to M (i.e. $N\phi \leftrightarrow \neg M \neg \phi$). Then, in view of c_2 above, we have a violation of intensional substitutivity:

$$(31) \quad cden(N(I = a), c_1) \neq cden(N(a = a), c_1).$$

²⁹See, e.g., Roberts (2015) for a brief overview and a sample account.

³⁰Stalnaker (1999) contains the classic treatment of presupposition in terms of a contextual-intensional semantic system. See Potts (2007) and Schlenker (2007) for treatments of expressives.

³¹I plug the context into the index slot, as usual.

In words, the sentence $N(I = a)$ in c_1 expresses the (always false) proposition that the speaker of every context is Trump, whereas $N(a = a)$ in c_1 expresses the (always true) proposition that Trump is Trump. Intensional substitutivity (content compositionality) is annulled the minute we define an operator in terms of quantification over the context parameter. This is completely analogous to how intensional operators destroy extensionality. We see how little these systems differ from one another on the foundational level, though they are meant to capture disparate notions.

Given our assumptions, we see that the prohibition on monsters coincides with content compositionality. This yields a test for whether a proposed formalism is free of monsters or not:

Content Compositionality Criterion. *If a language is content-compositional, then it is free of monsters.*³²

It is therefore enough to argue that a semantic system preserves content compositionality in order to say that it is an adequate direct reference system.

3 Towards Foundational Pragmatics

We saw that Kaplan's system is foundationally indistinguishable from extensional semantics. We also saw that it fails Kaplan's adequacy criterion. This co-occurrence is no coincidence. It is precisely because contexts are treated as run-of-the-mill objects that monsters are definable. In this section I'll give the rudiments of a formal system that gives the right interpretations of indexicals, but which treats the notions underlying indexicality in a foundationally appropriate way. The result, I will show, is content-compositional, and hence free of monsters. The guiding thought is that the notions of agency, context of use and direct reference should inform the very logic, and not just the theory, of indexicals.

3.1 Diagnosis

As noted, the definition of the monstrous operator M above is completely analogous to the definition of the temporal operator A (always) in §2.2. I repeat the two definitions for convenience:

$$(32) \quad \text{iden}(A(\phi), i) = 1 \text{ iff } \forall i' (\text{iden}(\phi, i') = 1).$$

$$(33) \quad \text{cden}(M(\phi), i, c) = 1 \text{ iff } \exists c' (\text{cden}(\phi, i, c') = 1).$$

³²This is an epistemic criterion like, say, the Tarski-Vaught test, and not an adequacy condition like Convention T.

In both cases there is an implicit *external* quantifier, which is common to all meaning clauses, making them general enough to serve as definitions. And in both cases there is an explicit *internal* quantifier (the “shifting” quantifier). The only difference is that in the indexical case the objects over which the internal quantifier ranges are contexts, and not indices. But this is not a difference in the logic, only in the informal gloss. Monsters are thus exactly on a par with intensional operators; except that intensional operators were the very reason for the development of intensional semantics, whereas monsters in indexical semantics are a sign that we went wrong. If we want to fix indexical semantics, we need to incorporate the philosophical difference between indices and contexts in the very framework.

Informally, indices are glossed as circumstances of *evaluation*, contexts as circumstances of *use*. The significance of this difference is brought out by considering the following sentence:

(34) It is raining.

This sentence can be *evaluated* with respect to any place and time, even ones in which no English speakers, or even no humans, have ever existed; but it cannot be *used* in such times and places. Kaplan, and others following him, have responded to this by identifying the contexts with a subset of the indices, namely with speaker-world-time-place quadruples such that the speaker is present at the world-time-place junction.³³ This captures the insight that contexts are somehow centered around or dependent on linguistic agents, but it does this somewhat superficially. In particular, it doesn’t answer to the active or dynamic aspect of language *use*, in contrast to the passive or static idea of *evaluation* of a sentence relative to a set of circumstances.

³³Such quadruples are called *proper contexts*. Kaplan (1989, 509) introduces them in order to capture the fact that sentences such as

(a) I am here now,

are true whenever uttered. For some reason he calls such sentences “logically true”. But in his logic, they only become true in virtue of a special clause tailored for this purpose (p. 544, clause 10), and not in virtue of the structure of his model theory, as we would expect logical truths to be. In any case, the ad-hoc clause is not sufficient, since there are sentences true whenever uttered which are not true in every proper context, e.g.:

(b) I am uttering a sentence.

Of course, we can further restrict the proper contexts by adding further ad-hoc clauses to handle each new “logically true” sentence. But this is not always forthcoming, for consider:

(c) I am uttering an English sentence.

This sentence will not be true in contexts in which, say, a French sentence is uttered, but *will* be true whenever uttered. This isn’t just a matter of including a language parameter in contexts. Here’s a sentence true whenever uttered, but true *only* in contexts in which *it* is uttered:

(d) I am uttering the result of prepending to its own quotation the string “I am uttering the result of prepending to its own quotation the string”.

The notion of proper context is not so easy to characterize.

Logically, the problem is that contexts are, on this treatment, a collection of self-standing objects, which is why they can be quantified over in the usual ways.³⁴ We need a way to “refer” to contexts which doesn’t carry ontological commitment, by somehow implementing the idea that language use, or linguistic agency, is that which creates the context.

Here’s a different angle on the same problem. For present purposes, linguistic agency can be identified with the capacity for *content generation*. The word “generation” means something like “bringing into existence”, and the implication is that there is a certain division of labor between the standing linguistic meanings on the one hand, and the context-situated linguistic agent on the other. On Kaplan’s system, content generation is modeled by abstracting the index parameter in the ternary denotation function $cden(e, i, c)$, yielding a binary function $cden(e, c)$ that returns the content of expression e in context c . If we further abstract the context parameter, we get a function $cden(e)$ from an expression to its character.

But from a foundational perspective, this is inadequate. Since contents and characters are modeled in complete analogy, the difference between the key terms – “generated” for contents and “standing” for characters – is idle in the formalism. To say that something is *generated in* context, we mean it is *unavailable out of* context. Characters are modeled as mathematical functions, and functions, on the usual extensional understanding (and Kaplan says nothing to the contrary), are such as to imply the existence of their arguments and values. It follows that the values of characters exist (are available) wherever characters exist, which means independently of context and agent. But the values of characters are just contents. So logically, contents are “standing” just as much as characters. Kaplan’s system doesn’t capture content generation in a foundationally sound sense of “generation”. This is another aspect of the previous problem, about contexts being modeled as just another kind of logically standard object. Here the problem is with the extensional interpretation of *function*. Indeed, Kaplan (1989) labels this move a matter of convenience.³⁵ But this convenience is what let the monsters in.

We have isolated two places that call for a revision of the formal system. First, we need a conception of contexts that doesn’t make them into full-blown objects, or we will not be able to rule out the internal quantifier of (33). This new conception should not rule out the external quantifier, however. So, more specifically, what we need is something like a new quantifier, which is universal in force but that can’t be nested. Second, we need to explicate characters, or standing rules of content generation, otherwise than in terms of a mathematical function understood extensionally.

³⁴By the object-blindness of extensional semantics, see §2.1.

³⁵P.505: “Just as it was convenient to represent contents by functions from possible circumstances to extensions (Carnap’s intentions [sic]), so it is convenient to represent characters by functions from possible contexts to contents.”

3.2 Prescription

Strawson (1950) famously expressed a skepticism regarding the possibility of formal semantics, especially in the case of context-sensitive words. He writes (327, original emphasis):

To give the meaning of an expression . . . is to give *general directions* for its use to refer to . . . particular objects . . . to give the meaning of a sentence is to give *general directions* for its use in making true and false assertions.

Stanley (2007, 4) judges Strawson’s view of meaning as a rule of use to be “vague and mystical”, and he praises Kaplan’s approach for overcoming this view through its appeal to mathematical functions from contexts to contents. By the considerations of the previous section, Stanley is wrong in his verdict. The notion of a mathematical function does not have the right properties for capturing standing meanings, and we might want to revisit Strawson’s vague and mystical “directions for use”, to see whether we cannot make something more scientific out of it.

What the functional representation of rules of use fails to capture is their force or mood. On the modern extensional view of mathematical functions, a function is only notationally different from a predicate. Every function is really just a concealed indicative statement of a correlation between several kinds of objects.³⁶ This view fails to capture one central feature of the intuitive, or at least historical, notion of function, which is reflected in the etymology of the word, namely the idea of *performing* an action or calculation.³⁷ Taking our cue from Strawson, we will not think of characters in terms of extensional mathematical functions, but as directives, instructions for a prospective linguistic agent to generate an item of content at a given context. In other words, we should think of characters, not in terms of indicative sentences asserting a static correlation, but in terms of imperatives.

The imperative is often presented as part of a paradigm, together with the indicative and the interrogative:

(35) Open the window!

(36) You open the window.

(37) Do you open the window?³⁸

Compared to their paradigm-mates, what characterizes imperatives is non-truth-aptness and second-person directedness.³⁹ The latter is exhibited in the fact that, although all three sentences are second-person in their orientation, only (35)

³⁶I’m referring to the fact that every expression of the form $f(x)$ can be systematically replaced by a predication $P(x, y)$, with suitably placed existential quantifier binding y .

³⁷From the Latin *fungi* – perform, execute, discharge. See <https://www.etymonline.com/>

³⁸See Kaufmann (2012, 2) for such a paradigm. The present simple tense sounds less natural in the plain indicative and interrogative, but we can ignore that.

³⁹See, e.g., Aikhenvald (2010, 3), Portner (2016), Roberts (2015).

achieves this without an overt second-person indexical. In addition (35) is not truth-conditional, which means that truth conditions cannot constitute its meaning.⁴⁰ There is, however, some relation between the imperative (35) and the indicative (36). Clearly, it is not entailment: From an utterance of (35), however felicitous, there is no way for me to infer (36). This is what led some writers to think of imperatives in terms of deontic modality, accounting for the second-person orientation in terms of a covert indexical. On such accounts, sentence (35) is equivalent to

(38) You must open the window,

the “must” being read deontically.⁴¹ But such accounts are at best descriptive. Foundationally speaking, indexicality and a non-alethic modality do not capture the imperative mood. The indexical “you” merely refers to the addressee; the imperative does something more: it *appeals* to them, targeting their capacity to act, their agency, to the exclusion of any other properties they might have. In addition, it is not just that the imperative fails to entail the indicative, as in deontic logic. The indicative *will follow* from the imperative, provided the addressee *complies* with it. We may say that the indicative *follows* from the imperative, *pending* the compliance of the addressee. It is therefore more precise to say that imperatives are characterized by being *appeals to agency*, and by their alethicity being *agent-pending*.

Let’s look at one special kind of imperative, a creation imperative such as:

(39) Bake a cake!

The corresponding indicative is:

(40) You bake a cake.

Putting considerations of baking time aside, the indicative (40) entails an existential sentence:

(41) There is a cake.

However, since in (39) a cake’s creation is being urged, it must be that that particular cake does not exist, at least not until the imperative is complied with. Thus (39) refers to a non-existent cake, or rather, to an agent-pending cake. Once the imperative is complied with by a suitable addressee, we have (40), and consequently (41). In this way the labor of ontological commitment, so to speak, is divided between the imperative and the complying agent. This is the general structure of what we need for content generation. I will refer to this mode of reference as *abstract reference*.

⁴⁰(36) doesn’t have a truth-value either, but its semantics, arguably, involves the truth conditions of (36) in a straightforward way.

⁴¹See, e.g., Isac (2015, 42ff) for an argument for the modal account.

3.3 Treatment

We want to revise Kaplanian semantics according to the considerations of the previous two subsections. The main locus of revision will be the metalanguage, and not the semantic theory. In this paper I’ll provide a brief sketch. A fuller development can be found in my (unpublished).

Let $M_{@}$ be a standard extensional first-order language, with the addition of a single term “@”, called a *pseudo-variable*. This term will refer to contexts in the abstract agent-pending mode described above.⁴² Formally, it corresponds to a universal quantifier which can’t be nested, but it isn’t interpreted, like the usual universal quantifier, in terms of the truth or satisfaction of every one of its instances. Sentences which contain it are understood not in the indicative mood, and therefore not truth-conditionally, but as imperatives, given a context, to interpret language in a certain way.

Given a context, we apply the context-feature functions with the pseudo-variable to get pseudo-singular-terms, e.g. *speaker(@)*, *time(@)*, that refer, in the abstract mode, to the speaker and to time of the context. They correspond to Strawson’s “*general directions* for [a term’s] use to refer to . . . particular objects . . .”. With them we can formulate pseudo-sentences, e.g.:

(42) *speaker(@) is president.*

These only have truth conditions when used in context, in other words, their truth-conditions are *agent-pending*. They correspond to Strawson’s “*general directions* for [a sentence’s] use in making true and false assertions”. We let $M_{@}$ quantify in the standard sense (non-abstractly) over intensional indices. This lets us formulate terms referring to pseudo-contents:

(43) $\{i : \textit{speaker}(@) \textit{ is president at } i\}$.

(43) is a pseudo-proposition: an agent-pending set of indices. Philosophically, it models an instruction for an agent to generate a proposition.

This is, in brief outline, the *abstract pragmatic metalanguage* $M_{@}$. In it we formulate a semantic theory which, like Kaplan’s *cden*(e, i, c), will yield the desired positive results, but without giving rise to monsters. Let the object-language $L_{@}$ be syntactically identical to L_c of §2.3. We’ll show that the sentence *PI* (“I am president”) expresses the same content as the sentence *Pa* (“Trump is president”) in contexts in which Trump is the speaker, and only in those. First, note that *Pa* expresses the same proposition regardless of context. We use an intensional (non-indexical) denotation function to express this fact:

⁴²There are some complications that I will gloss over for the sake of brevity. For example, the context doesn’t “come into existence” in the same way that the cake of the example does. The context is the collection of circumstances of an utterance, and as such, it doesn’t literally depend on the agent for its existence. But it does turn from a mere static self-standing “index of evaluation”, to a dynamic “context of use” in virtue of the speaker’s linguistic agency. I expand on this in Kashtan (unpublished).

$$(44) \quad iden(Pa) = \{i : Trump \text{ is president in } i\}.$$

This is as before. The situation changes when indexicals come on the scene. A sentence containing indexicals doesn't have a content outside of context, it has its content "generated" in context. On the Kaplanian system, this is modeled by a function:

$$(45) \quad cden(PI) = \lambda c. \{i : speaker(c) \text{ is president in } i\}.$$

But this gives rise to monsters. To fix this, we replace the variable c with the pseudo-variable $@$, to get:

$$(46) \quad iden_{@}(PI) = \{i : speaker(@) \text{ is president at } i\}.$$
⁴³

$iden_{@}(e)$ can be thought of as a function from expressions to pseudo-contents, which are directives for generating proper contents in context. In a particular concrete context c_1 (where Trump is speaker and president), the proper contents get "generated" by taking a pseudo-content and replacing the pseudo-singular-terms in it with the actual features of c_1 . Thus, the term " $speaker(@)$ " is replaced with "Trump", etc.⁴⁴:

$$(47) \quad iden_{c_1}(PI) = \{i : Trump \text{ is president at } i\}.$$

If in c_2 Hillary Clinton is the speaker, then we have:

$$(48) \quad iden_{c_2}(PI) = \{i : Hillary \text{ Clinton is president at } i\}.$$

Since " Pa " has the same content in every context, we have:

$$(49) \quad iden_{c_1}(Pa) = \{i : Trump \text{ is president at } i\}.$$

The sentence Pa therefore has the same content as the sentence PI in the context c_1 , but not in the context c_2 . This is as desired.

3.4 Checkup

It remains to show that $iden_{@}(e)$ doesn't give rise to monsters. In order to see this, consider first that the function $iden_{@}(e)$ is definable along the lines of $iden(e)$. The non-indexicals get exactly the same treatment in $iden_{@}(e)$ as they do in $iden(e)$, e.g.:

$$(50) \quad iden_{@}(a) = \lambda i. Trump.$$

⁴³I subscripted the $@$ since it is not a genuine variable. The resulting denotation function then has the signature of $iden(e)$, and not of $cden(e, c)$, a fact I marked in the spelling.

⁴⁴And not with $speaker(c_1)$, since " c_1 " is not an available term, either in $M_{@}$ nor in $L_{@}$. See Kashtan (unpublished).

The (pseudo-)semantic values of indexical expressions are given in terms of the context features and the pseudo-variable:

$$(51) \quad iden_{@}(I) = \lambda i. speaker(@).$$

We call on the fact that $iden(e)$ is definable in a single metalinguistic sentence, which we have abbreviated $def_i(base_{L_i}, D^{L_i}, D^I)$ taking the lexicon $base_{L_i}$ and the domain D^{L_i} as parameters.⁴⁵ In a similar fashion we can express the lexicon of $iden_{@}(e)$ in a single sentence $base_{L_{@}}$. We assume the object domain to be context-dependent as well, and denote it $D^{@}$. Then we can mimic the definition of $iden(e)$, to get a definition $def_i(base_{L_{@}}, D^{@}, D^I)$.

That no monsters are definable here follows from two observations we’ve made earlier. In §2.3 we stated that any content-compositional system will be monster free; and in §2.2 we saw that a semantic function built along the lines of $iden(e)$ will be content-compositional. At a given context, say c_1 , the contents expressible will be exactly the values of $iden_{c_1}(e)$. We get a definition of $iden_{c_1}(e)$ if we take the definition of $iden_{@}(e)$ and replace every pseudo-content by the content it receives in c_1 . This definition has the form of the standard definition of $iden(e)$, which means that $iden_{c_1}(e)$ is content-compositional. It follows by the content-compositionality test that it is monster-free.

What does the job is the fact that the pseudo-variable @, though it is a kind of universal quantifier, cannot be nested. Philosophically this is motivated by the fact that contexts are not self-standing objects, and therefore that there is no self-standing domain of them over which a regular quantifier can range. Contexts only become real relative to a situated agent, and only such an agent can refer to objects in terms of their positions in context. The general rules for such reference, on the other hand, must abstract from any particular context in order to be general. This is the division of labor, between the abstract but general “rules of use” and the situated agent, which makes indexical reference possible.

4 Conclusion

My aim in this paper has been to suggest an alternative approach to the question of monstrous operators and their prohibition. Kaplan’s contribution to the semantics of indexicals is often said to be the disentanglement of the indexical aspect of language from the intensional aspect. But Kaplan only went half-way, and his conceptual account of indexicality remained intertwined with his understanding of intensionality. In addition, he was not sufficiently clear about whether his theory is proposed as an empirical theory about natural languages, or as a contribution to foundational semiotics. By separating sharply the indexical from the intensional,

⁴⁵See the item labeled “Definability” in §2.2.

and the foundational from the empirical, I hope to have shown that the prohibition on monsters touches on very fundamental questions in the abstract study of language.

In closing, I'd like to mention a couple of applications to which we can put the directive semantics sketched in §3. The first is the notion of schema in the formal sciences. Important theories such as Peano arithmetic or Zermelo-Fraenkel set theory irreducibly rely on schemata for their axiomatization. For example, the principle of arithmetical induction is axiomatized in first-order Peano arithmetic through the infinitely many instances of the induction axiom schema (where ϕ is replaced by a formula with one free variable):

$$(52) \quad \phi 0 \wedge \forall n(\phi n \rightarrow \phi(n + 1)) \rightarrow \forall n \phi n.$$

The schema itself is not part of the language of Peano arithmetic, so in an important sense, since we can't use infinitely many sentences, the theory of arithmetic is not *given* in the language of arithmetic. One way to give the theory in a single sentence is with the truth predicate.⁴⁶ By Tarski's theorem, however, this can only be done in a distinct, and stronger, metalanguage. This might not be a problem in the case of the language of arithmetic, as there are many languages strong enough to define truth for it. But in the case of set-theory, the idea of a stronger metalanguage leads to certain philosophical difficulties.

The directive semantics of §3 offers a different approach. Rather than stating an infinite collection of sentences, we think of a schema as a directive for generating sentences to be taken as axioms. In this way we are not expressing the theoretical content of the schema, the totality of its instances, but its practical content, the rule of use which directs us in making new instances of it according to need. Unlike a semantic metalanguage, the language $M_{@}$ is not stronger than its object-language since it can't define truth for it.⁴⁷

The second application is to the problem of the concept of truth. As mentioned just now, Tarski's theorem says that the truth predicate for a language can only be defined in a strictly stronger metalanguage. Many philosophers have been uncomfortable with the fact that this leads to an unbounded hierarchy of metalanguages, and have tried to circumvent Tarski's result in various ways. In fact, the hierarchy is not only unintuitive, it is on the face of it untenable. For the mere assertion of its existence assumes a language which quantifies over absolutely all languages, including itself. But since languages are individuated by their interpretation, and quantification involves individuation, it follows that this ultimate metalanguage will contain an interpretation, i.e. a truth theory, for every language; including itself. And this is again impossible by Tarski's theorem.⁴⁸ This I consider to be the hard problem of Tarskian semantics.

⁴⁶This is the truth predicate's best known use. See Picollo and Schindler (2019) for a recent discussion.

⁴⁷More about this in Kashtan (unpublished). Compare this account of schemata with that of Field (2006).

⁴⁸See Priest (1984), Putnam (1990, 13f) for discussions of this problem.

The directive semantics of §3 gives us a handle on this problem. For it allows us to state things with generality but without ontological commitment. Using the pseudo-variable, with its universal non-nestable quantificational power, a directive metalanguage like $M_{@}$ might successfully serve as an overarching metalanguage in which we can speak about *all* languages (all indicative languages) without regress.⁴⁹ If this is correct, then the question of the prohibition on monsters reaches far deeper than Kaplan suspected.

References

- Aikhenvald, A. Y. (2010). *Imperatives and commands*. Oxford: Oxford University Press.
- Anand, P., & Nevins, A. (2004). Shifty operators in changing contexts. *Semantics and Linguistic Theory*, 14, 20–37.
- Atkin, A. (2005). Peirce on the index and indexical reference. *Transactions of the Charles S. Peirce Society*, 41(4), 161–188.
- Bar-Hillel, Y. (1954). Indexical expressions. *Mind*, 63(251), 359–379.
- Briciu, A. (2018). On context shifters and compositionality in natural languages. *Organon F: Medzinárodný Časopis Pre Analytickú Filozofiu*, 25(1), 2–20.
- Burgess, A., & Sherman, B. (Eds.). (2014). *Metasemantics: New essays on the foundations of meaning*. Oxford: Oxford University Press.
- Carnap, R. (1942). *Introduction to semantics*. Cambridge: Harvard University Press.
- Carnap, R. (1955). Meaning and synonymy in natural languages. *Philosophical Studies*, 6(3), 33–47.
- Cohen, J., & Michaelson, E. (2013). Indexicality and the answering machine paradox. *Philosophy Compass*, 8(6), 580–592.
- Donnellan, K. S. (1966). Reference and definite descriptions. *Philosophical Review*, 75(3), 281–304.
- Field, H. (2006). Compositional principles vs. schematic reasoning. *The Monist*, 89(1), 9–27.
- García-Carpintero, M. (1998). Indexicals as token-reflexives. *Mind*, 107(427), 529–564.
- Gupta, A. (2012). *Truth, meaning, experience*. Oxford: Oxford University Press.
- Heim, I., & Von Stechow, P. (Eds.). (2011). *Intensional semantics* (Lecture notes).
- Isac, D. (2015). *The morphosyntax of imperatives* (Oxford studies in theoretical linguistics) (Vol. 59). Oxford: Oxford University Press.
- Israel, D., & Perry, J. (1996). Where Monsters Dwell. In J. Seligman & D. Westerståhl (Eds.), *Logic, language and computation* (Vol. 1, pp. 1–303). Stanford: Csl Publications.
- Kaplan, D. (1978). Dthat. In P. Cole (Ed.), *Syntax and semantics* (Vol. 9, pp. 221–243). San Diego & London: Academic Press.
- Kaplan, D. (1989). Demonstratives: An essay on the semantics, logic, metaphysics and epistemology of demonstratives and other indexicals. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan* (pp. 481–563). Oxford: Oxford University Press.
- Kaufmann, M. (2012). *Interpreting Imperatives*. Springer.
- McCullagh, M. (2017). Kinds of monsters and kinds of compositionality. *Analysis*, 78(4), 657–666.
- Montague, R. (1970). Pragmatics and Intensional logic. *Synthese*, 22(1–2), 68–94.
- Piccolo, L., & Schindler, T. (2019). Deflationism and the function of truth. *Philosophical Perspectives*. <https://doi.org/10.1111/phpe.12113>.

⁴⁹This is the topic of Kashtan (unpublished).

- Portner, P. (2016). Imperatives. In M. Aloni & P. E. Dekker (Eds.), *The Cambridge handbook of formal semantics* (Cambridge handbooks in language and linguistics) (pp. 593–626). Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9781139236157.021>.
- Potts, C. (2007). The expressive dimension. *Theoretical Linguistics*, 33(2), 165–198.
- Predelli, S. (2014). Kaplan's three monsters. *Analysis*, 74(3), 389–393.
- Priest, G. (1984). Semantic Closure. *Studia Logica*, 43(1–2), 117–129.
- Putnam, H. (1990). *Realism with a human face*. Cambridge, MA: Harvard University Press.
- Rabern, B. (2013). Monsters in Kaplan's logic of demonstratives. *Philosophical Studies*, 164(2), 393–404.
- Rabern, B., & Ball, D. (2019). Monsters and the theoretical role of context. *Philosophy and Phenomenological Research*, 98(2), 392–416.
- Roberts, C. (2015). *Conditional plans and imperatives: A semantics and pragmatics for imperative mood*. In: Proceedings of the 20th Amsterdam Colloquium.
- Roberts, L. D. (1994). The foundations of Kaplan's direct reference theory for demonstratives. *Philosophia*, 23(1–4), 91–116.
- Russell, B. (1905). On denoting. *Mind*, 14(56), 479–493.
- Schlenker, P. (2002). A plea for monsters. *Linguistics and Philosophy*, 26(1), 29–120.
- Schlenker, P. (2007). *Expressive presuppositions*. <https://doi.org/10.1515/TL.2007.017>.
- Schlenker, P. (2011). Indexicality and de se reports. In M. von Stechow & M. G. Portner (Eds.), *Handbook of semantics*. Oxford: Oxford University Press.
- Stalnaker, R. (2014). *Context*. Oxford: Oxford University Press.
- Stalnaker, R. C. (1999). *Context and content: Essays on intentionality in speech and thought*. Oxford: Oxford University Press.
- Stanley, J. (2007). *Language in context: Selected essays*. Oxford: Clarendon Press.
- Stanley, J., & Gendler Szabó, Z. (2000). On quantifier domain restriction. *Mind and Language*, 15(2–3), 219–261.
- Strawson, P. F. (1950). On referring. *Mind*, 59(235), 320–344.
- Tarski, A. (1956). The Concept of Truth in Formalized Languages. in *Logic, Semantics, Metamathematics*. Hackett.
- Von Stechow, A. (2003). Feature deletion under semantic binding: Tense, person, and mood under verbal quantifiers. *Proceedings-NELS*, 33, 379–404.
- Westerståhl, D. (2012). Compositionality in Kaplan-style semantics. In W. Markus, H. Wolfram, & E. Mächery (Eds.), *The Oxford handbook of compositionality*. Oxford: Oxford University Press.

Compositionality in Truth-Conditional Pragmatics



Adrian Briciu

Abstract In the past decade various linguists and philosophers (e.g. Pagin, Pelletier, Recanati, Westerståhl, Lasersohn) have proposed a weakening of the standard interpretation of compositionality for propositional content. Their move is motivated by the desire to accommodate radical forms of context sensitivity within a systematic account of natural languages. In this paper I argue against weakening compositionality in the way proposed by them. I argue that weak compositionality fails to provide some of the expected benefits of compositionality. First, weak compositionality fails to provide systematic meaning-rules which can handle forms of context-sensitivity that are not amenable to explanation in terms of a fixed and limited set of contextual parameters. Secondly, I argue that weak-compositionality fails to play any role in explaining speakers' ability to calculate the semantic values of complex expressions. I conclude that weak compositionality is not a viable alternative to standard interpretations of compositionality, and that it doesn't offer an acceptable way to accommodate radical forms of context-sensitivity within a systematic account of natural languages. Given the central role that weak-compositionality plays in recent approaches to natural language (e.g. in truth-conditional pragmatics) this also casts doubt on the viability of these projects.

Keywords Compositionality · Context-sensitivity · Formal semantics · Truth-conditional pragmatics

1 Introduction: Meaning and Compositionality

Formal semantics aims to model the fundamental meaning properties and meaning relations of natural languages with the tools of formal logics. It aims to build formal characterizations of natural languages that can serve as explanatory models of our

A. Briciu (✉)

Department of Philosophy and Communication Studies, West University of Timisoara, Timisoara, Romania

semantic competence – of what a speaker knows in knowing her language. The formal theory can start from the basic insight that to know the meaning of a sentence is to know the conditions under which the sentence is true, and the conditions under which the sentence is false. Then as a way of specifying our competence with meaningful natural language expressions, the theory will seek to systematically pair each meaningful sentence of a language with its truth-conditions. The theory is descriptively adequate if its pairings of sentences with truth-conditions, match native speakers' intuitions about the conditions under which these sentences are true and the conditions under which they are false. Moreover, given that speakers of natural languages have the ability to produce and understand complex meaningful expressions that they have never encountered before, one central goal is to explain how linguistic competence extends to novel expressions. Since the most plausible explanation is that speakers derive their competence with novel sentences from their competence with the constituents of the sentences and competence with ways in which expressions can combine, a semantic theory should derive in a finite number of steps the truth-conditions of sentences from the meanings of their simple parts and their syntactic structures. For this reason, the vast majority of semanticists believe that compositionality is a fundamental property of natural languages: that the meaning of complex expressions is a function of the meaning of their immediate constituents and syntactic structure.

Still, philosophers and linguists with a Wittgensteinian bent believe that the existence of pervasive forms of context-sensitivity in natural languages, and the apparent unruliness of language use, threatens the very project of formal semantics.¹ They believe that natural languages exhibit forms of context-sensitivity that cannot be treated in terms of a fixed set of contextual parameters (in the way in which expressions like “I” or “that” are treated) and that this is incompatible with the principle of compositionality. More recently, though, a motley coalition of linguists and philosophers deny that there is incompatibility between (some versions of) compositionality and radical forms of context sensitivity.² The position defended by the latter is the focus of this paper. I will argue that their proposals to weaken the principle of compositionality, in order to accommodate within a compositional framework recalcitrant data, loses the theoretical benefits promised by compositionality.

The plan of the paper is the following. This first section presents a framework in which claims about meaning and compositionality can be clearly formulated. Section 2 presents the challenge from radical forms of context-sensitivity to the very project of formal semantics, and Sect. 3 presents the truth-conditional pragmatics' (henceforth TCP) proposal to accommodate radical forms of context-sensitivity within a compositional framework. The last two sections are dedicated to an

¹Ziff (1972), Searle (1978), Travis (1978, 1997), Margalit (1979), Moravcsik (1994), Bezuidenhout (2002), Carston (2002), Recanati (2004)

²See Pelletier (2003), Pagin (2005), Pagin and Pelletier (2007), Westerståhl (2012), Lasersohn (2012) and references therein.

extensive criticism of the TCP approach and its proposed interpretation of the principle of compositionality.

1.1 Linguistic Meaning and Propositional Content

Following Kaplan (1989) and Lewis (1980) it is customary to distinguish two types of meaning: what an expression means independently of any context of utterance (or *linguistic meaning*) and what an expression means relative to a context of utterance (or *propositional content*).

The linguistic meaning of an expression is the convention associated with that expression and something like a rule of use: it tells what an expression can say when used in any arbitrary context. For example, the linguistic meaning of the first person pronoun “I” can be thought of as a rule that states that “I” when used in an arbitrary context refers to the speaker of that context. Linguistic meaning can be modeled as a function from possible contexts of utterance to propositional contents. The linguistic meaning of some expressions is a non-constant function that returns different contents for different contexts of utterance (e.g. indexicals like “I”, “today”, descriptions like “the tallest man in town”, etc.), while the linguistic meaning of others is a constant function that returns the same content at any context of utterance (proper names like “David Kaplan”, numerals like “two”, etc.). Contexts, as formal objects of the theory, have the job to represent the concrete situation in which language use takes place. Since any use of an expression is done by an agent, at a place and time in a possible world, contexts can be represented as a sequence of individuals consisting of an *agent*, a *time*, a *location* and a *world*. An essential assumption in semantic theorizing is that we can give the linguistic meaning of any natural language expression in terms of a limited and fixed set of contextual parameters, even if we *may* need to extend the list of parameters beyond agent, time, location and world.

The propositional content of an expression can be thought of as the information that determines the denotation of an expression at any possible state of affairs. Propositional content can be modeled as a function from circumstances of evaluation (which at minimum are possible worlds but, in principle, could be richer) to denotations: individuals for singular terms, sets for predicates, truth-values for sentences. The idea is, in somewhat simplified terms, that the linguistic meaning of a sentence determines a unique content with respect to any given context, and the content determines, in its turn, a unique denotation with respect to any given circumstance of evaluation.

A fundamental tenet that underlies all semantic theorizing is that for any context of utterance the linguistic meaning of a sentence determines its truth-conditions at that context, and that all context sensitivity can be handled in terms of a fixed and limited set of contextual parameters, more or less along the lines in which expressions like “I”, “today”, “here” are treated. The idea is that the theory assigns a finite number of meanings to simple expressions and uses a finite number of

rules to derive, out of these meanings, the truth-conditions of every sentence of that language. As I will detail in the second section, this general picture has been vigorously contested by authors who are skeptical towards the very possibility of formal semantics.

1.2 Varieties of Compositionality

Depending on whether we require that it is the linguistic meaning of complexes or their propositional content that is a function of the values of constituents we obtain different versions of compositionality. Here is how semanticists traditionally formulated compositionality for linguistic meaning and for propositional content, respectively.³

A. *Compositionality of linguistic meaning*: The linguistic meaning of a complex expression is a function of the linguistic meaning of its constituents and of its syntactic structure. More precisely, a semantics \mathbf{I}^* is linguistic meaning compositional (LM compositional) iff for any syntactic rule α there is function f such that for any two expressions e_i, e_j if $\alpha(e_i, e_j)$ is meaningful then:

$$\mathbf{I}^*(\alpha(e_i, e_j)) = f(\alpha, \mathbf{I}^*(e_i), \mathbf{I}^*(e_j)).$$

A semantics fails to be LM compositional if for some expressions e_i, e_j, e_n , and syntactic rule α ,

$$\mathbf{I}^*(e_j) = \mathbf{I}^*(e_n) \text{ and } \mathbf{I}^*(\alpha(e_i, e_j)) \neq \mathbf{I}^*(\alpha(e_i, e_n)).$$

In plain words, a semantics fails to be LM compositional if substitution of synonyms is not meaning preserving in that semantics.

Since content is assigned to expression-context pairs, in order to formulate compositionality for content we need to take into account the role that context plays in the determination of the content of complexes.

B. *Strong compositionality for content*: the content of a complex expression relative to a context is a function of the content of its constituents at that context and of its syntactic structure. More precisely, a semantics \mathbf{I}^* is *strongly compositional*

³See Kaplan (1989, 507) where both varieties are given informally. For their formal rendering see Pagin and Westerståhl (2010, 259–260), Dever (2006, 634), Szabó (2010, 258–260). Given that linguistic meaning is a property of expressions themselves, linguistic meaning will be assigned directly to expressions, and given that propositional content is a property of expressions at contexts, content will be assigned to expression-context pairs. Furthermore, since linguistic meaning is a function from contexts to propositional content, and propositional content is a function from circumstances to extensions, a semantics \mathbf{I}^* which assigns linguistic meaning directly to expressions is the curried version of a semantics \mathbf{I} which assigns propositional content to expression-context pairs. That is, for any expression e and any context C , $\mathbf{I}(e, C) = \mathbf{I}^*(e)(C)$.

iff for every syntactic rule α there is a function f such that for any two expressions e_i, e_j and for any context C if $\alpha(e_i, e_j)$ is meaningful at C then:

$$\mathbf{I}(\alpha(e_i, e_j), C) = f(\alpha, \mathbf{I}(e_i, C), \mathbf{I}(e_j, C))$$

A semantics fails to be strongly compositional if for some expressions e_i, e_j , syntactic rule α and some contexts C_1, C_2

$$\begin{aligned} \mathbf{I}(e_i, C_1) = \mathbf{I}(e_i, C_2) \text{ and } \mathbf{I}(e_j, C_1) = \mathbf{I}(e_j, C_2) \\ \text{and } \mathbf{I}(\alpha(e_i, e_j), C_1) \neq \mathbf{I}(\alpha(e_i, e_j), C_2). \end{aligned}$$

In other words, a semantics fails to be strongly compositional if a complex expression varies its content across contexts of utterance but its constituents have unvarying contents across the very same contexts of utterance. A relevant consequence of strong compositionality is that the content of a complex expression depends on the context only in so far as the contents of its constituents do. If the content of a complex expression is context dependent this should be traceable to the context-dependency of at least one of its simple constituents.

Recently, various theorists have argued that context should be given a more substantive role in the determination of the content of complexes, and that this is compatible with the spirit of compositionality. We get, then, another principle of compositionality for content:

C. Weak compositionality for content: The content of a complex expression relative to a context C is a function of the contents that its constituents have at C and of C itself. More precisely: a semantics \mathbf{I} is weakly compositional iff: for every syntactic rule α there is a function f such that for any expressions e_i, e_j and for any context C if $\alpha(e_i, e_j)$ is meaningful then

$$\mathbf{I}(\alpha(e_i, e_j), C) = f(\alpha, \mathbf{I}(e_i, C), \mathbf{I}(e_j, C), C).$$

A semantics fails to be weakly compositional if for some expressions e_i, e_j, e_n , syntactic rule α and context C : $\mathbf{I}(e_j, C) = \mathbf{I}(e_n, C)$ and $\mathbf{I}(\alpha(e_i, e_j), C) \neq \mathbf{I}(\alpha(e_i, e_n), C)$.

Observe that according to weak compositionality (as opposed to strong compositionality) context is taken as an extra argument of the composition function and *the contribution that this extra argument makes can be non-vacuous*. A consequence of weak compositionality is that the context sensitivity of complex expressions need not be traceable to the context-sensitivity of some of its constituents. Weak compositionality allows that context determines the content of a complex expression in ways that go over and above determining the contents of its constituents.⁴

⁴Strong compositionality is a proper generalization of weak compositionality, which, in its turn, is a proper generalization of linguistic meaning compositionality. For proofs see Westerståhl (2012). For an alternative proof and a further discussion of how weak and strong compositionality for content interact with various types of context-sensitivity see Briciu (2018).

2 Radical Context-Sensitivity and Semantic Underdetermination

Ever since its inception skeptical voices claimed that the very project of formal semantics is wrongheaded and doomed to fail. According to them, because English exhibits forms of context-sensitivity that cannot be handled simply in terms of the lexical and syntactic properties of its expressions, formal semantic theories cannot give correct truth-value predictions even for simple English sentences like (1).

(1) The leaves are green

The following scenario, imagined by Travis (1997, 89–90), is meant to show this.⁵ Suppose Pia paints the leaves of a Japanese russet maple tree green for a photographic installation. Upon ending the job she might utter truly (1) while pointing to the leaves. Later, a botanist friend seeking green leaves for an experiment on green-leaves chemistry drops by. Pia offers her botanist friend the leaves she has just painted, and utters (1) while pointing towards them. But now, she might for all the paint, utter falsely (1) while pointing to the leaves. Intuitively, the two utterances of (1) have different truth values; the first is true while the second is false, although the brute state of the leaves did not change in between the two utterances. According to Travis, (1) is neither ambiguous nor elliptical; it contains neither vague nor indexical expressions, nor are our intuitions about its truth value the result of what might be indirectly conveyed by its respective utterances (i.e. intuitions about the truth or falsity of what is implicated by those utterances).⁶

If skeptics are right, whether (1) is true or false at the imagined scenario depends not only on its linguistic meaning and how the world is, but on a multitude of potentially unrepeatable and formally intractable factors, like the participants' immediate interests, purposes and concerns. Allegedly, the difference in truth-values is due to the fact that (1) has different truth-conditions at the two contexts of utterance: it is true in the photographer context iff the leaves *appear green* at the time of the utterance; while it is true in the botanist context iff the leaves *are naturally green* at the time of the utterance. Skeptics further argue that the difference in truth-

⁵According to skeptics, radical context-sensitivity affects virtually any natural language sentence. Arguments similar with Travis' have been put forward concerning rather pedestrian sentences like (2) "It is raining" (Recanati 2002), or (3) "The cat is on the mat" (Searle 1978), or (4) "The ham sandwich stinks" (Recanati 2010), or (5) "The snow is white" (Moravcsik 1994). This skepticism is also shared by linguists like Chomsky (2002). For discussions of many more such examples see Bezuidenhout (2002), Recanati (2004), Cappelen and Lepore (2005), Szabó (2007), García-Carpintero (2006).

⁶Obviously some of these claims were contested. For example Kennedy and McNally (2010) argue that (1) is ambiguous because color terms are ambiguous between gradable and non-gradable interpretations. Given that there are many other similar arguments put forward by skeptics that do not involve color terms and that TCP accepts, for the sake of the argument I will go along with skeptics and truth-conditional pragmatists and accept their claim that no vagueness, ellipsis or ambiguity is involved in (1).

conditions of (1) is not determined by its linguistic meaning. Their argument can be reconstructed in the following way:

- (A) The two utterances of (1) have different truth values. (*Data*)
- (B) The difference in truth value is due to the fact that (1) expresses different contents at the two contexts of utterance.
- (C) The constituents of (1) do not vary their content across the relevant contexts of utterance
- (D) Since the linguistic meaning of constituents of (1) is not context-sensitive, the linguistic meaning of (1) should determine the very same content at the two context of utterance.

From (B) and (D) it follows:

- (E) Underdetermination: the contextual variation in the content of (1) is not determined by its linguistic meaning (*Conclusion*)⁷

According to Travis, as far as the meanings of constituents of (1) and its syntax go, whether we can predicate truly “green” of the leaves in question is an open matter: on some occasions we can, on others we can’t. What determines whether an utterance of (1) is true in the photographer context and false in the botanist context is not just its linguistic meaning and the brute state of the leaves, but also an intricate web of immediate interests, intentions and beliefs of the conversational partners.

Skeptics believe that for virtually any natural language sentence what that sentence literally means, together with formal aspects of context (who is speaking when and where), plays some role in determining its truth-conditions, but not an exhaustive one. Its truth-conditions depend on factors that cannot be made completely explicit in the semantic analysis of the sentence. One reason is that these factors are not fixed: there is *no constant set of factors* that determines the truth conditions of a sentence relative to *any* context of utterance. Another reason is that the list of factors relevant for determining the truth-conditions of a sentence is open-ended: “information from virtually anywhere and about virtually anything might have a bearing”⁸ on truth-conditions. Even for the simplest sentences, human interests, concerns and beliefs can play a role in determining whether they are true or false at a given context of use, and there is no determinate boundary at the outset on which facts could turn out to be relevant for the interpretation of a sentence. Skeptics take this to show that no systematic account of the meaning properties of natural languages, with the tools of formal logic, is possible. If true, then an important part of our linguistic competence might lie beyond the reach of systematic theorizing.

⁷Needless to say, defenders of formal semantics try to resist the above argument by rejecting some of its premises. Borg (2004a, b) and Cappelen and Lepore (2005) deny that the data put forward by skeptics are semantically relevant, Predelli (2005) denies premise (B) and argues that (1) expresses the same content at the two contexts where the difference in truth-value is the result of evaluating the content for truth at different circumstances; Szabó (2001) and Rotschild and Segal (2009) deny premise (C) and argue that “is green” is context-sensitive after all.

⁸Carston (2002, 2). This is also the central argument in Searle (1978), Bezuidenhout (2002) and Recanati (2004, chapter 9).

Underdetermination of truth-conditions by linguistic meaning also entails failure of strong compositionality: the content of sentences (relative to a context) *is not* determined by the content of their constituents (at that context) and the way the constituents are syntactically combined. Premise (B) together with (C) entail the following claim:

- (F) Some complex expressions –e.g. (1) – vary their content across contexts of utterance although the content of their constituents remains stable across the very same contexts.

Arguments surrounding the truth of these premises constitute the bulk of the debate on whether the project of formal semantics is wrongheaded or not. Given that TCP accepts these arguments, I will not judge their worthiness, nor will I try to adjudicate on whether natural languages exhibit radical forms of context sensitivity. Rather, my aim is to discuss if the TCP proposal to accommodate radical context-sensitivity within a weakly compositional account delivers the explanatory benefits that we expect from compositional theories of natural languages.

3 Truth-Conditional Pragmatics and Weak Compositionality

TCP accepts that natural languages exhibit forms of context-sensitivity which cannot be treated by fixing the values of a limited set of contextual parameters, and accepts that this brings about semantic underdetermination. But it claims that a compositional account of natural languages can still be given, although not as initially conceived. According to TCP, semantics and pragmatics mix in determining truth-conditional content: pragmatic factors (i.e. factors not mandated by the lexical and syntactic properties of expressions) play a role in the determination of contents of sentences (at contexts of use). This is where weak compositionality comes in. A theory that allows for pragmatic intrusion through and through fails to be strongly compositional, but it *can* be weakly compositional. By making use of weak compositionality, TCP promises to model natural languages by systematically pairing sentences with their truth-conditions (i.e. what formal semantics traditionally aims to do) in a way that can accommodate recalcitrant cases brought up by skeptics like Travis: sentences whose truth-conditions depend on a potentially open-ended number of pragmatic factors.⁹ Here is how Recanati (2010, 127) summarizes the main idea behind TCP:

[T]ruth conditional pragmatics is the view that the effects of context on the content need not be traceable to the linguistic material in the uttered sentence. Some effects of context on content are due to the linguistic material (e.g. the context sensitive words or morphemes which trigger the search for contextual values), but others result from “top down” pragmatic processes that take place not because the linguistic material demands it, but because

⁹Pagin and Pelletier (2007, 32) are explicit about this.

utterance's content is not faithfully or wholly encoded in the uttered sentence, whose meaning requires adjustment or elaboration in order to determine an admissible content for the speaker's utterance.

In TCP pragmatics kicks in not only to derive what is conveyed by an utterance of a sentence (e.g. to derive the conversational implicatures), but plays a role also in determining the truth-conditions of that sentence.¹⁰ Although both semantic processes (like indexical resolution) and primary pragmatic processes determine semantic content, they are distinct in that the first, but not the latter, are required by the lexical and/or syntactic properties of expressions. Because they are lexically and/or syntactically required, semantic operations are *mandatory* (in the sense that in their absence a sentence fails to express a truth-evaluable content) while pragmatic operations, are merely *optional* (in the sense that in their absence a sentence might still express a truth-evaluable content).

To get the gist of TCP, consider how it analyses (4).

(4) The ham sandwich stinks

There are numerous contexts in which (4) expresses exactly what its linguistic meaning says, namely that the contextually salient ham sandwich stinks. For example, if one sorts rotten food from good one, and utters it, (4) is taken to be true iff the salient ham sandwich stinks. But imagine that in order to maximize speed and efficiency restaurant workers tend to refer to their customers by the dish that they order. If used in such a context (4) is true just if the person who ordered the ham sandwich stinks; its truth conditions at this context involve a person and not a ham sandwich.

TCP believes that these intuitions about (4) constitute bona-fide data that theories which seek to model competence with natural language meanings must account for.¹¹ It puts forward the following proposal. The linguistic meaning of each simple constituent of (4) determines together with the context of utterance the literal content of that expression at that context. The contents thus determined are combined step-by-step following the syntactic structure of (4) into the content of complex constituents ending with the content of (4) at that context. At some contexts though (e.g. the restaurant context) the contribution of the noun-phrase to the content of (4) does not involve a sandwich, and thus it is not what is determined by its lexical and syntactic properties. Its contribution is, in part, determined by pragmatic operations. Loosely speaking we could say that there is a context-specific pragmatic function that maps dishes into their orderers which is relevant for the interpretation of (4)

¹⁰For this purpose Recanati (2004, 23–37) distinguishes two types of pragmatic operations: *primary* (they play a role in the determination of truth-conditions) and *secondary* (they play a role solely in the derivation of conversational implicatures).

¹¹This phenomenon was first discussed in Nurnberg (1995). Of course, the first-blush reaction that defenders of formal semantics have in the face of these examples is to deny their semantic significance: to deny that intuitions about metonymic uses of (4) are to be treated on a par with those of literal use, and that a common treatment of both is desirable. For a discussion along these lines see Stanley (2007, 206–207).

at the restaurant-context. What results from combining the content of the parts determined by their linguistic meaning according to the syntactic structure of (4) is only an intermediate stop in the overall process of determining the content of (4). Relative to some contexts, its content is determined in part by context-specific pragmatic operations.

The term used by TCP for these types of pragmatic operations is *modulation*. Formally, modulation can be represented as a function that takes us from the content determined by the lexical and syntactic properties of an expression to a conversationally more appropriate content – that is, a function from content to content. Modulation is *context-specific* in that it can vary with the context of utterance: for every expression-context pair $\langle e, C \rangle$ there can be a distinct modulation function $mod(e, C)$ that determines the content of that expression at that context.¹²

For example, the modulation function that determines the content of the compound nominal in (4) at the restaurant-context is distinct from the one that determines its content in the context of sorting food.¹³ According to TCP, at some contexts C , what a simple expression e contributes to the content of complexes, is not the content determined solely by its lexical and syntactic properties $\mathbf{I}(e, C)$, but a pragmatically determined content $mod(e, C)(\mathbf{I}(e, C))$. Furthermore, pragmatic functions can also operate on complexes directly, as is the case with the compound nominal in (4). The content of a complex itself can be the result of a contextually salient pragmatic function:

$$Mod(\mathbf{I}(\alpha(e_i, e_j), C)) = mod(\alpha(e_i, e_j), C) \\ \left(f\left(\alpha, mod(e_i, C)(\mathbf{I}(e_i, C)), mod(e_j, C)(\mathbf{I}(e_j, C))\right) \right)$$

It is clear that pragmatic functions which operate on complexes destroy strong compositionality. For example, such an account of (4) fails to satisfy strong compositionality since it allows (4) to vary its content across contexts of utterance although its simple constituents keep constant contents across the very same contexts. Nevertheless such an account of (4) *can* satisfy weak compositionality.

¹²As Recanati puts it, “modulation itself is context-sensitive: whether or not modulation comes into play, and if it does, which modulation takes place, is a matter of context” (Recanati 2010, 19). In their formal apparatus both Pagin and Pelletier (2007) and Recanati (2010) make use of a general modulation function *mod* which sole purpose is to determine the particular, context-specific modulation functions: *mod* takes pairs of expressions e and contexts C as arguments and delivers, for each such pair, the contextually appropriate modulation function $mod(e, C)$.

¹³Within this account literalness can be treated as a limiting case: the context-specific function that delivers the content of “the ham sandwich” in the context of sorting food is the identity function.

4 Against Weakening Compositionality

If the threat to the project of formal semantics is that truth-conditions of natural language sentences vary in an un-systematic way, finding out that natural languages satisfy weak compositionality won't save the project of building a systematic theory of their meaning properties. In this section I'll argue for this, taking into account two different ways in which TCP can be implemented.

Truth-conditional pragmatics can have two distinct, but formally equivalent, architectures. In one, interpretation takes place in one fell swoop employing context-sensitive meaning-rules: rules that introduce meaning-operations which vary with the context of utterance. In the other, interpretation is a two-steps process whereby context-insensitive meaning rules operate on the lexical and syntactic properties of expressions, and they feed context-specific pragmatic functions which determine the truth-conditions of sentences at contexts of utterance.

4.1 Weakly Compositional Meaning Rules

In order to keep things clear let me rehearse, first, the distinction between *rules* and *operations*, a distinction that applies both at the level of syntax and at the level of semantics. *Syntactic rules* state how expressions of a language combine to form larger grammatical expressions. Here is an example of syntactic rule: if e_1 is an expression of category NP and e_2 is an expression of category VP then concatenating e_1 and e_2 , in this order, results in an expression of category S. This particular rule introduces one type of *syntactic operation* by which expressions combine, namely concatenation. *Meaning-rules* state how the meaning of complex expressions with a certain syntactic structure is obtained. Here is an example of meaning-rule: if $e_1 \wedge e_2$ is a complex expression formed by concatenating e_1 and e_2 , in this order, and the meaning of e_1 is a function whose domain contains the meaning of e_2 then the meaning of $e_1 \wedge e_2$ is the value of the meaning of e_1 for the meaning of e_2 as an argument: $\mathbf{I}^*(e_1 \wedge e_2) = \mathbf{I}^*(e_2)(\mathbf{I}^*(e_1))$. This meaning-rule introduces one type of *meaning-operation* by which meanings combine, namely functional application^{14,15}. Importantly, this rule specifies the semantics \mathbf{I}^* in a way that allows to derive the semantic value of complex expressions once we match

¹⁴This is but one of many compositional rules available to theorists. Other rules can introduce other types of operations for various complex expressions. For a discussion see Chung and Ladusaw (2004, 2–14)

¹⁵In a sense, meaning-rules interpret syntactic ones. Each syntactic rule states that expressions of certain syntactic categories can combine to form expressions of a certain syntactic category, and determine the operation by which they combine. And each meaning-rule states how (i.e. by which operation) the meanings of complex expressions with a certain syntactic structure are built from the meanings of their constituents.

up the variables in the rule with the constituents that correspond to them in each particular expression.

What type of meaning rules are needed by TCP? More precisely, what type of rules are needed by a theory which allows for pragmatics to determine what content an expression has relative to a given context of use? If constituents do not vary their content, what is, then, the source of this variation? It must be the manner in which the contents of constituents combine at different contexts. Then the theory needs *context-sensitive meaning-rules*: rules that introduce meaning-operations which vary with the context of utterance. Such meaning-rules look handy for TCP, for they promise to help explain how (1) and (4) vary their content across contexts of utterance in the absence of a corresponding variation in the content of their simple constituents. Furthermore, such rules are weakly compositional.

A context-sensitive meaning rule looks the following way: if $e_1 \wedge e_2$ is a complex expression formed by concatenating e_1 and e_2 , in this order, then for any context C , there is an operation O such that the semantic value of $e_1 \wedge e_2$ at C is the result of combining by O the semantic value of e_1 at C with the semantic value of e_2 at C . This, though, is rather uninformative; obviously, at every context the semantic values of constituents combine in one way or another. For context-sensitive rules to be of any use to natural language semantics they must specify *how* semantic values combine: they must specify, for any context, the value of the variable O for that context.

To get a flavor of how context-sensitive rules work, consider a toy language which has the same vocabulary and syntax as English but its meaning-rules introduce operations which vary across contexts of utterance as a function of the height of the speaker. Its meaning-rule states that if $e_1 \wedge e_2$ is a complex expression formed by concatenating e_1 and e_2 , in this order, then for any C , the content of $e_1 \wedge e_2$ at C is the result of applying the content of e_1 at C to the content of e_2 at C if the speaker of C is shorter than 1.60 m, *or* the content of $e_1 \wedge e_2$ at C is the result of applying the content of e_1 at C to the content of e_2 at C and applying negation to the content thus obtained, if the speaker of C is taller than 1.60 m. In this language the sentence “John walks” expresses the proposition that John walks if uttered by a speaker shorter than 1.60 m and expresses the proposition that John doesn’t walk if uttered by a speaker taller than 1.60.¹⁶ Its context-sensitive rule specifies how operations vary with the context of utterance.

¹⁶It is obvious that this rule is not strongly compositional. But it is weakly-compositional. Under the assumption that a fragment of English, of which Vertical English is an extension, is weakly compositional, it can be shown that Vertical English is weakly compositional too. If the initial language is weakly compositional then extending it with the above rule does not destroy its weak compositionality. In Vertical English for any two sentences $e_1 \wedge e_2$ and $e_1 \wedge e_3$ and any context C , if $a_c \leq 1.60$ m and $I(e_2, C) = I(e_3, C)$ then $I((e_1 \wedge e_2), C) = I((e_1 \wedge e_3), C)$ – the content of constituents combine through functional application. And for any two sentences $e_1 \wedge e_2$ and $e_1 \wedge e_3$ and any context C if $a_c > 1.60$ m and $I(e_2, C) = I(e_3, C)$ then $I((e_1 \wedge e_2), C) = I((e_1 \wedge e_3), C)$ – the content of constituents combine through the complex operation described. Thus, the failure condition of weak compositionality, given in Sect. 1.2, does not obtain.

What kind of context-sensitive rules are suitable for TCP? According to TCP, sentences vary their contents across contexts of use in virtue of various pragmatic factors, so the rules it needs must introduce meaning-operations which vary with pragmatic factors. More precisely, the weakly compositional rules needed by TCP must systematically match meaning-operations with the corresponding pragmatic factors in terms of which they vary. For example, TCP needs rules of the following form: if $e_1 \wedge e_2$ is a complex expression formed by concatenating e_1 and e_2 , in that order, then for any C if the conversational partners have the intention i and concern b at C then the content of $e_1 \wedge e_2$ at C is the result of combining the content of its constituents by operation X (say, functional application), *or* if the conversational partners have intention i and concern d at C the content of $e_1 \wedge e_2$ at C is the result of combining the content of its constituents by operation Y (say, predicate restriction).

But this is problematic for anyone wedded to the idea that natural languages exhibit radical forms of context sensitivity; that is, the type of context sensitivity that can't be handled by fixing the value of some definite contextual parameter(s). If there is no determinate boundary at the outset on which factors can turn out to be relevant in determining the set of truth-conditions of a sentence, it is impossible to pair each meaning-operation with those in terms of which it varies. To pair them, theorists must predict ahead of time each and all of the potentially open-ended number of pragmatic factors that, in principle, can be relevant for the interpretation of sentences with a certain syntactic structure. In other words, the type of rules that systematically pair meaning operations (i.e. ways of combining meanings) with pragmatic factors are incompatible with the very idea that the truth-conditions of a sentence depend on an indefinite number of pragmatic factors i.e. factors which cannot be fully encoded into the sentence meaning.¹⁷

To put this worry from a different angle: if natural languages are weakly compositional and their expressions are radically context sensitive, a single syntactic structure will contribute in more than one way to the interpretation of complex expressions, and its contribution will vary freely with the context of utterance. A weakly compositional meaning rule, then, will have to specify for the syntactic structure it interprets several manners of combining meaning. If these ways of combining meanings vary across contexts not as a function of a fixed and limited set of parameters, but vary together with a potentially open-ended number of highly specialized and intricate arrangements of intentions, interests, and expectations of the conversational partners, then stating such rules is highly problematic, for several reasons.

First, nothing short of a full model of human practical reasoning must be packed into meaning-rules. Even theorists, who are optimistic about achieving such a model, should find this a serious drawback for a theory that aims to model the meaning properties of natural languages. Secondly, if meaning rules are sensitive to the vicissitudes and peculiarities of each possible context of use, it is not clear at all why TCP needs compositionality. Compositionality is desirable because

¹⁷See Recanati (2004, 194), Travis (1996, 451), Bezuidenhout (2002, 105)

it delivers rules which are an effective procedure for calculating the semantic values of complexes such that executing the procedure requires no imagination or cleverness, but is a matter of merely following instructions carefully. But if truth-conditions depend on the intentions, practical concerns and the common assumptions of the conversational partners, then calculating them cannot be a matter of merely following instructions carefully. Rather, deriving them requires assigning mental states to conversational partners and reconstructing their practical reasoning. In other words, derivation of truth-conditions becomes an inference to the best explanation.

Thirdly, stating such rules requires doing precisely what the underdetermination claim denies it can be done: to predict ahead of time each and all of the potentially open-ended number of factors that might turn out to be relevant for the interpretation of an expression. And TCP accepts, and is motivated by, the alleged existence of semantic underdetermination.

There are several defensive moves that TCP can try, in order to alleviate these worries, but none of them are convincing.

As a first defensive move TCP might point out the obvious fact that not every possible way of combining meaning is permissible in English (nor in any other natural language). Although, in principle, there are an open-ended number of propositions that (1) can express, it can't express *any* proposition. There are limits on how much speakers can tinker with the meaning of sentences they use. TCP will propose, then, to distinguish those operations on semantic values which are permissible in English from those which are not.¹⁸ Unfortunately this won't help. Even if we assume that there are a small number of meaning operations, given that there are a potentially open-ended number of pragmatic factors with which these operations vary, a theory must give indefinitely many pairings between them.¹⁹ Obviously, introducing the pairs by listing them is not an option. The meaning-rule must be, or include, a systematic procedure that matches every meaning-operation with the appropriate pragmatic factors. But, again, this is incompatible with the main point of semantic underdetermination, namely that the linguistic meaning of a sentence is essentially open-ended: no set of rules can determine its truth-conditions for all possible contexts of utterance.²⁰

A different defensive move for TCP is to use meaning-rules that introduce operations which vary with a fixed and limited number of parameters. This,

¹⁸Recanati (2010, 11) and Pagin and Pelletier (2007, 57) hint towards this move

¹⁹That TCP can do with a small number of meaning operations already concedes a lot. It looks to me that TCP is committed to the claim that there are a potentially open-ended number of meaning-operations. This follows directly from two of its other claims: (a) that a sentence can, in principle, express an open-ended number of propositions, each particular to a given context, and (b) that this variation need not be traceable to a corresponding variation in the content of the simple constituents, but that it can be the result of combining the content of constituents by different operations at different contexts

²⁰See Searle (1978) and Margalit (1979).

though, involves giving up on the idea that natural languages exhibit radical forms of context sensitivity, the very phenomenon that motivated the appeal to weak-compositionality in the first place. Why is this solution incompatible with radical forms of context-sensitivity? If natural languages make use of this type of rules, any sentence would vary its content only as a function of a fixed and limited number of parameters, precisely what is denied by those who believe that natural languages exhibit radical forms of context sensitivity. In other words, such rules are of no use for TCP.²¹

Finally, TCP might argue that there is no need for meaning-rules to pair each meaning-operation with the pragmatic factors in terms of which they vary. They might point to analyses of demonstratives which rely on the notion of demonstratum or salience, without giving an explanation of how it is determined what is demonstrated, or what is salient. In fact, Pagin and Pelletier (2007, 58–59) suggest that determining what meaning-operation is at work in a context is similar to selecting the referent of a demonstrative relative to a context. We don't have a general and fully satisfactory theory that will tell us how to predict what is the most salient person, object or relation in a certain context of utterance, but we don't take this limitation to impinge on the systematicity of semantic theories. We should take the same attitude when it comes to operations on meaning. Then TCP should be content with formulating very general rules of the form: if e_3 is a complex expression and e_1 and e_2 are its immediate constituents, then the content of e_3 at any given context C , is the result of combining the content of e_1 at C with the content of e_2 at C in the way relevant at C .

There are good reasons to believe that this move is not available to TCP and that appeal to reference resolution for demonstratives is not helpful. In fixing the reference of a demonstrative relative to a context, the determination of the saliency profile of the context is beyond the reach of semantics. So is the determination of the salience profile of the context when it comes to determining which way of combining meanings is relevant at that context of utterance. A formal theory will tell us what "That is red" means relative to a context, but it won't tell us why "that" refers to x and not to y , relative to that context, other than that x and not y is the salient (or demonstrated) object. What object is the most salient one (or the demonstrated one) in a context is beyond the reach of semantic theories. There is no reason to suppose that semantic theories should tell us what particular object satisfies the property of being the most salient object (or the demonstrated object) at a given context. Thus a semantic theory might deliver an analysis of "That is red" along the following lines: if the speaker of "that is red" refers with the utterance of

²¹This is acknowledged also by Lasersohn. He writes with respect to such rules: "the contextual effects that threaten compositionality are of a much more thorough-going nature than the effect illustrated in [this rule], and do not lend themselves to an analogous treatment" (2012, 186).

“that” therein to x and to nothing else, then this sentence, as uttered in this context, is true if and only if x is red.²²

Can we really say that determining the way in which meanings combine (relative to a context) is beyond the scope of semantics, just as determining the most salient (or the demonstrated) object of a context, is beyond the scope of semantics? This seems absurd, for just as one can’t have a theory about the combinatorial and structural properties of expressions (i.e. syntax) without an account of how expressions combine, one can’t have a theory about the meaning properties of a language in without an account of how meanings themselves combine.

In other words, such a rule is useless for a theory that seeks to model the fundamental meaning properties of a natural language. For such a rule doesn’t tell us *how* to calculate the semantic value of complex expressions, since it doesn’t introduce any meaning operation. And stating how to calculate the semantic values of complex expressions is precisely what semantic rules are expected to do. To say that the semantic values of complexes (with a given syntactic structure) are the result of combining the semantic values of their constituents in the relevant way is to say something trivial.²³ It is part and parcel of any theory that models the meaning-properties of a language to assign meaning to simple expressions *and* to determine the semantic effects of combining those meanings in given syntactic configurations. A theory that employs the type of rule described above won’t tell us what the semantic effects of combing expressions in a certain syntactic configuration are.

4.2 Context-Specific Pragmatic Functions

These problems persist even if one prefers a different architecture for TCP, one in which derivation of truth-conditional content is a two-step process. For example, in the first step the linguistic meaning of each simple constituent of (1) determines, together with the context of utterance, a propositional content for that expression at that context. At this step all meaning-rules are context-insensitive. Relative to any context of utterance, the contents of simples (as determined by their lexical properties) are combined through *functional application* into the literal content of

²²When it comes to reference fixing this is a strategy advocated, among others, by Borg (2004c, 2012), Higginbotham (1989), and Heck (2014).

²³Moreover, there is another reason to doubt that TCP can successfully appeal to theories of demonstratives that rely on salience in order to make a case for rules which do not introduce meaning operations. Even if the explanation of how an object becomes salient within a context of utterance is beyond the scope of a theory of meaning, there is a substantive story to be told about this. But there is no substantive story to be told about how one meaning-operation becomes more salient than another one. To say that one way of combining meaning is more salient than another is just to say that one interpretation of a complex expression is more salient, or more readily available to than another one.

(1) at that context. In the second step, for each context of utterance a *context-specific pragmatic function* takes the propositional content determined by the lexical and syntactic properties of (1) and delivers the truth-conditional content of (1) at that context.²⁴

If there are contextual ingredients in the truth-conditions of (1) which are provided through pragmatic functions, a theory that aims to predict for every context under what conditions (1) is true at that context must be able to predict for every context the right pragmatic function. More generally, if for any sentence *S* and context *C* there is a pragmatic function that determines the truth-conditions of *S* at *C*, and it is possible that for each sentence-context pair there is a distinct function, then TCP must provide a systematic procedure to calculate the right pragmatic function, without making use of independent knowledge of the truth-conditions of *S* at *C*. In the absence of this there is no systematic way to derive the truth-conditions of sentences and the threat posed by radical context-sensitivity remains unaddressed.

In order to be explanatorily rewarding, that is, in order to be able to derive truth-conditions for sentences (at contexts of utterance), TCP needs to generalize in a substantive way over these particular pragmatic functions: it needs to state a rule, or a finite set of rules, which determine for each context the correct pragmatic function. This is somewhat problematic, since selecting the correct pragmatic function, out of a potentially open-ended number of such functions, does not seem to be a matter of following rules, but one of recognizing intentions and of reasoning through inference to the best explanation. In TCP, explaining how a given sentence comes to have the truth-conditions that it has, is partly an intentional explanation: it involves attributing certain intentions and practical concerns to the conversational partners. Then the assignments of truth-conditions that TCP makes are always defeasible, for the simple reason that intentional explanations are always defeasible: they can always be overridden if enough new evidence is adduced to account for the subject’s linguistic behavior. In fact, Recanati points out that “a distinguishing characteristic of pragmatic interpretation is its defeasibility, [the fact] there is no limit to the amount of contextual information that can affect the interpretation” (Recanati 2004, 54).

The problem for TCP is not that finding the right pragmatic function is necessarily impossible. The problem is that finding the right pragmatic function is, essentially, an intentional explanation: it requires assigning intentions, beliefs

²⁴In fact, this is closer to the organization of TCP that Recanati (2010) and Pagin and Pelletier (2007) work with. It is easy to see that the two ways of organizing TCP are formally equivalent. In the two-step version, the content of an expression $\alpha(e_i, e_j)$ at a context *C* is determined by a context-specific pragmatic function $mod(\alpha(e_i, e_j), C)$ which takes as argument what is determined by the lexical and syntactic properties:

$$\mathbf{I}(\alpha(e_i, e_j), C) = mod(\alpha(e_i, e_j), C)(f(\alpha, (\mathbf{I}(e_i, C), \mathbf{I}(e_j, C))).$$

Notice that this is formally equivalent with $\mathbf{I}(\alpha(e_i, e_j), C) = f \circ mod(\alpha(e_i, e_j), C)(\alpha, (\mathbf{I}(e_i, C), \mathbf{I}(e_j, C)))$

where $f \circ mod(\alpha(e_i, e_j), C)$ is a complex function obtained by combining the composition function *f* and the modulation function $mod(\alpha(e_i, e_j), C)$. This corresponds to the way of building TCP where a context-specific meaning operation combines the content of constituents into the content of the complex in one fell swoop.

and practical reasoning to conversational partners, and as such is always defeasible. Thus, we might legitimately doubt that TCP's answer to the skeptic's challenge really provides substantive progress towards a systematic theory. The skeptics told us that some truth-value predictions made by formal theories are bound to be wrong, while TCP tells us that any of its predictions might turn out to be wrong, for any such prediction relies on abductive reasoning.

TCP might try any of the defensive strategies discussed in the previous section, but none of them will work. I'll consider, again, the last of them. TCP might argue that there is no need to specify a procedure that determines, for each context of utterance, the particular pragmatic function that delivers truth-conditions at that context of utterance. Again, TCP might point to analyses of demonstratives which rely on the notion of demonstratum (or of salience), which do not determine what object is demonstrated (or is more salient) at the context of utterance. TCP would claim is that for each context there is a pragmatic function at work, but it is beyond the scope of the theory to determine for each context what that function is. This move is not satisfactory, for TCP will give truth-conditions to (1) of the following form:

"The leaves are green" is true at a context C iff the objects that the speaker intends to refer to with that utterance of "the leaves" satisfy the property that the speaker intends to predicate about them with that utterance of "are green".

This amounts to saying that (1) is true at C iff the proposition that the speaker intended to express by (1) at C is true at the circumstances determined by C.²⁵ But this is wholly uninformative and it satisfies the aim of pairing sentences with their truth conditions in an extremely shallow way.

Finally, irrespective of what architecture TCP prefers, if the derivation of truth-conditions essentially involves attributing intentions, beliefs and practical reasoning to conversational partners, it is not at all clear why TCP needs compositionality. Again, compositionality is desirable because it promises to deliver rules which are an effective procedure for calculating the semantic values of complex expressions such that executing the procedure requires no imagination or cleverness, but is a matter of merely following instructions carefully. But, obviously, this is not the case with the attribution intentions and beliefs or the reconstruction of practical reasoning. Letting the composition function take context as an extra argument might not be against the letter of compositionality, but it is against its spirit.

²⁵TCP might point out that there are limits on what propositions a sentence can express, because there are limits on how much one can tinker with the meaning of sentences: even if a sentence can express indefinitely many propositions, it can express any proposition. This, though, doesn't make its analysis of (1) more informative.

5 Conclusions

It is fair to conclude that weak compositionality offers no way out from the skeptical challenge concerning the possibility of a systematic semantics of natural languages. True, weak compositionality shows that, formally, there is no incompatibility between radical context-sensitivity and *some* version of compositional interpretation of complex expressions. But this is far from being enough. We're interested in compositionality because of the explanatory benefits it promises to bring. I argued that we have good reasons to doubt that weak-compositionality delivers these benefits, if we accept the existence of radical forms of context-sensitivity.

One such benefit is that compositionality ensures that there are rules by which theories can derive the truth-conditions of natural language sentences from the meaning of simple expressions. But if one is convinced that natural language expressions are radically context-sensitive, then weak compositionality won't help with this. Acceptance of radical context-sensitivity amounts to accepting that there are elements in the content of sentences that are not determined by the lexical and syntactic properties of sentences themselves, but are provided through pragmatic functions. Unless we are given a procedure that matches at the outset each possible context of utterance with its associated pragmatic function, weakly compositional theories can't systematically derive truth-conditions. As yet, no theory has provided such procedure. I suggested, moreover, that for someone who accepts underdetermination, as TCP theorists do, there are good reasons to doubt that such rules or procedures can be given, for they require that the theorist be able to tell ahead of time each and all of the facts that might turn out relevant for the interpretation of a sentence.

Given the central position that weak-compositionality occupies in truth-conditional pragmatics these arguments cast doubt over the viability of the entire project. If weak-compositionality fails to provide the explanatory benefits that we expect from compositionality, it is doubtful that TCP can deliver on its advertising claim, namely doing what formal semantic theories aimed but allegedly failed to do: offer a systematic account of our linguistic competence. This, of course, is not to say that radical contextualists, like Searle (1978) or Travis (1997), are right. A systematic account of natural languages might still be possible. But if one accepts that there are contextual ingredients in the truth-conditions of sentences which are provided through free pragmatic functions (i.e. are not linguistically mandated), weak-compositionality is not going to help in providing such an account.

Acknowledgments Thanks to Emma Borg, Manuel García Carpintero, Zoltán Gendler Szabó, Max Kölbel, Josep Macià, Sara Packalén, Peter Pagin, Dan Zeman and to the audience at the first *Context, Cognition and Communication Conference* and the University of Warsaw for extremely helpful questions, comments and objections. I acknowledge the financial support received from the Spanish Ministry of Economy and Competitiveness, EXCELENCIA program, project no. FFI2016-80636-P: *Foundations and Methods of Natural Language Semantics*.

References

- Bezuidenhout, A. (2002). Truth-conditional pragmatics. *Philosophical Perspectives*, 16, 105–134.
- Borg, E. (2004a). Formal semantics and intentional states. *Analysis*, 64(3), 215–223.
- Borg, E. (2004b). *Minimal semantics*. Oxford: Clarendon Press.
- Borg, E. (2004c). The place of referential intentions in linguistic content. *Manuscripto* 32(1), 85–122.
- Borg, E. (2012). *Pursuing meaning*. Oxford: Oxford University Press.
- Briciu, A. (2018). On context shifters and compositionality in natural languages. *Organon F*, 25(1), 2–20.
- Cappelen, H., & Lepore, E. (2005). *Insensitive semantics*. Oxford: Blackwell.
- Carston, R. (2002). *Thoughts and utterances*. Oxford: Blackwell.
- Chomsky, N. (2002). Explaining language use. In *New horizons in the study of language and mind*. Cambridge: Cambridge University Press.
- Chung, S., & Ladusaw, W. (2004). *Restriction and saturation*. Cambridge, MA: MIT Press.
- Dever, J. (2006). Compositionality. In E. Lepore & B. Smith (Eds.), *The Oxford handbook of philosophy of language* (pp. 633–666). Oxford: Oxford University Press.
- García-Carpintero, M. (2006). Recanati on the semantics/pragmatics distinction. *Critica*, 38(112), 35–68.
- Heck, R. (2014). Semantics and context-dependence: Towards a Strawsonian account. In B. Sherman & A. Burgess (Eds.), *Metasemantics: New essays on the foundations of meaning* (pp. 327–364). Oxford: Oxford University Press.
- Higginbotham, J. (1989). Elucidations of meaning. *Linguistics and Philosophy*, 12(4), 465–517.
- Kaplan, D. (1989). Demonstratives: An essay on the logic, metaphysics, and epistemology of demonstratives and other indexicals. In J. Perry & H. K. Wettstein (Eds.), *Themes from Kaplan* (pp. 481–614). Oxford: Oxford University Press.
- Kennedy, C., & McNally, L. (2010). Color, context, and compositionality. *Synthese*, 174(1), 79–98.
- Lasnik, P. (2012). Contextualism and compositionality. *Linguistics and Philosophy*, 35(2), 171–189.
- Lewis, D. (1980). Index, context, and content. In S. Kanger & S. Öhman (Eds.), *Philosophy and grammar* (pp. 79–100). Dordrecht: Reidel.
- Margalit, A. (1979). Open texture. In A. Margalit (Ed.), *Meaning and use* (pp. 141–152). Dordrecht/Hingham: D. Reidel.
- Moravcsik, J. (1994). Is snow white? In P. Humphreys (Ed.), *Patrick Suppes: Scientific philosopher: Volume 3* (pp. 71–87). Dordrecht: Kluwer Academic Publishers.
- Nurnberg, G. (1995). Transfers of meaning. *Journal of Semantics*, 12(2), 109–132.
- Pagin, P. (2005). Compositionality and context. In G. Preyer & G. Peter (Eds.), *Contextualism in philosophy: Knowledge, meaning and truth* (pp. 303–348). Oxford: Oxford University Press.
- Pagin, P., & Pelletier, J. (2007). Context, content and communication. In G. Preyer & G. Peter (Eds.), *Context-sensitivity and semantic minimalism* (pp. 25–62). Oxford: Clarendon Press.
- Pagin, P., & Westerståhl, D. (2010). Compositionality I: Definitions and variants. *Philosophy Compass*, 5(3), 250–264.
- Pelletier, F. J. (2003). Context-dependence and compositionality. *Mind & Language*, 18, 148–161.
- Predelli, S. (2005). *Contexts. Meaning, truth, and the use of language*. Oxford: Oxford University Press.
- Recanati, F. (2002). Unarticulated constituents. *Linguistics and Philosophy* 25(3), 299–345.
- Recanati, F. (2004). *Literal meaning*. Cambridge: Cambridge University Press.
- Recanati, F. (2010). *Truth-conditional pragmatics*. Oxford: Oxford University Press.
- Rothschild, D., & Segal, G. (2009). Indexical predicates. *Mind & Language*, 24(4), 467–493.
- Searle, J. (1978). Literal meaning. *Erkenntnis*, 13, 207–224.
- Stanley, J. (2007). *Language in context: Selected essays*. Oxford: Oxford University Press.
- Szabó, Z. (2001). Adjectives in context. In I. Kenesei & R. Harnish (Eds.), *Perspectives on semantics, pragmatics, and discourse* (pp. 119–146). Amsterdam: John Benjamins.

- Szabó, Z. (2006). Sensitivity training. *Mind and Language* 21(1), 31–38.
- Szabó, Z. (2010). The determination of content. *Philosophical Studies*, 148(2), 253–227.
- Travis, C. (1978). Meaning versus truth. *Dialogue*, 17(03), 401–430.
- Travis, C. (1996). Meaning's role in truth. *Mind*, 105(419), 451–466.
- Travis, C. (1997). Pragmatics. In B. Hale & C. Wright (Eds.), *A companion to the philosophy of language* (pp. 87–107). Oxford: Blackwell.
- Westerståhl, D. (2012). Compositionality in Kaplan style semantics. In M. Werning, W. Hinzen, & E. Machery (Eds.), *The Oxford handbook of compositionality* (pp. 192–219). Oxford: Oxford University Press.
- Ziff, P. (1972). *Understanding understanding*. Ithaca: Cornell University Press.

Occasion-Sensitivity and What Is Said



Claudia Picazo

Abstract Although both contextualism and occasionalism hold an underdeterminacy thesis with respect to natural language, they disagree about whether other structured representational items (such as mental representations or structured propositions) exhibit a similar form of underdeterminacy. In particular, contextualism assumes that the properties or mental representations expressed by a use of a predicate are underdeterminacy-free. I argue that this view faces three worries. On the linguistic side, if there are underdeterminacy-free properties to be had, it is mysterious why we are unable to coin words corresponding to them (i.e., underdeterminacy-free words). On the cognitive side, it remains to be explained how we relate to these non-encodable properties. And on the metaphysical side, it is not clear whether our communicative intentions are precise enough to determine an underdeterminacy-free property.

Keywords Semantic underdeterminacy · Contextualism · Occasion-sensitivity · Truth-conditional pragmatics

1 Introduction

Since the last decades, pragmatic approaches to language gain ground. Truth-conditional semantics is often replaced with truth-conditional pragmatics. According to the latter, the truth-conditions of a token of a non-indexical sentence *S* are not exhausted by the linguistic meaning of *S*. Rather, truth-conditions need to be supplemented by contextual information.¹ There are, however, different versions of

¹To some philosophers ‘need to be’ might sound too strong. These philosophers might prefer ‘often are’.

C. Picazo (✉)
Department of Philosophy, University of Granada, Granada, Spain

truth-conditional pragmatics, and they model the content of speech acts in different ways. Hence exploring how to think about linguistic content in the age of pragmatics becomes an important task.

My aim here is to examine two different forms of truth-conditional pragmatics and provide some arguments tipping the scale towards the more radical one. Using Borg's terminology (2012), I will call these views contextualism and occasionalism. In particular, I will focus on what is usually called radical contextualism and in Searle and Travis' radical pragmatic views. Contextualism and occasionalism have much in common. Both these views include an underdeterminacy principle to the effect that the linguistic meaning of a non-indexical sentence *S* underdetermines the truth-conditions of *S*'s tokens. In this sense, we can say that they take truth-conditional content to be occasion-sensitive. However, there is a relevant difference between them. Whereas contextualism has it that there are classical propositions to be had, i.e., some structured representational items that are free of underdeterminacy (such as mental representations or structured propositions), occasionalism rejects the existence of representational items whose truth-conditions are independent of the occasion of use.

The contextualist picture I will attack is admittedly a simplification. It is not always clear what is meant by 'proposition' in debates about truth-conditional pragmatics, so it might very well be the case that some versions of contextualism are almost indistinguishable from occasionalism. However, the contextualist reading of the underdeterminacy principle—a reading according to which utterances of the same non-indexical sentence can express different propositions at different contexts—suggests a certain picture worth exploring. The picture is the following. Just like the word 'he' refers to different men at different contexts, words as 'green' refer to different properties at different contexts. Thus, the word 'green' has different satisfaction conditions at different contexts. By contrast, the properties that can be expressed with this word do not. They are the building blocks of propositions—the primary bearers of truth-conditions, insensitive to the context at which they are tokened.

The problem that I see in this view is that it is not at all clear what these properties amount to or how we manage to relate to them. They are not encodable in natural language. Most, if not all, sentences in natural language seem to suffer from underdeterminacy. This is something that the version of contextualism I will focus on admits. So the building blocks of propositions cannot be equivalent to the (underdetermined) properties we encode in natural language (the property of being green, the property of being superficially green, and so on). In the following sections I will argue that these ineffable properties introduce several problems. On the linguistic side, if these underdeterminacy-free properties are the referents of our words, it is mysterious why we are unable to coin words corresponding to them (i.e., underdeterminacy-free words). On the cognitive side, it remains to be explained how we relate to these ineffable properties. And on the metaphysical side, it is not clear that our communicative intentions are precise enough to determine an underdeterminacy-free property. Here is the plan. In Sect. 2 I briefly present a salient argument in support of the underdeterminacy principle—Travis cases. Section 3 is

about the difference between contextualism and occasionalism. In Sect. 4 I present several arguments against the tenability of contextualism. These are not knockdown arguments. Rather, they are intended to encourage those with sympathies towards truth-conditional pragmatics to explore more radical views as occasionalism. In Sect. 5 I sketch how an occasionalist approach could look like.

2 Travis Cases

According to some linguistic theories, the truth conditions of a large class of declarative sentences (those that contain no indexicals or vague terms) are determined by the linguistic meaning of the simple expressions in the sentence together with some rules of composition. Following this idea, semantic theories are conceived as theories that yield the truth-conditions of non-indexical declarative sentences on the basis of the meaning of the simple expressions in those sentences and the rules of composition. Let us call this project truth-conditional semantics.

A number of philosophers, including saliently Searle and Travis,² have argued that truth-conditional semantics is ill-conceived. In a nutshell, their idea is that truth-conditions depend on something besides linguistic meaning and syntax (even in absence of indexicals). Their arguments consist on a number of examples showing that the truth-conditions (or, more generally, satisfaction conditions) of some sentences are contingent upon an indefinite number of elements that cannot be made part of the theory, including background assumptions about the world and specific features of the situation of use. Although Searle and Travis don't always frame their views in the same way or are interested in the same things, they both hold an argument to the effect that semantics (linguistic meanings) does not determine satisfaction conditions. Here are some well-known examples:

Pia's Japanese maple is full of russet leaves. Believing that green is the colour of leaves, she paints them. Returning, she reports, 'That's better. The leaves are green now.' She speaks truth. A botanist friend then phones, seeking green leaves for a study of green-leaf chemistry. 'The leaves (on my tree) are green,' Pia says. 'You can have those.' But now Pia speaks falsehood. (Travis 1997, p. 89)

Consider the words, 'Hugo weighs 79 kilos' and the following situation: when Hugo steps on the scale in the morning, it reads 79 kilos, and that is a stable result. However, it is now after lunch; fully dressed (in winter clothing), Hugo would register 81 kilos on the scales. Now consider two speakings of the words. For the first, Hugo must weigh 79 kilos, and no more, to qualify for some sporting event. There is a discussion as to whether he does qualify. Odile, who has seen him step on the scale, tries to settle the matter by revealing her information on the subject. She says, 'Hugo weighs 79 kilos.' For the second, Hugo is about to step on to a very delicate trestle bridge across a ravine which can take a maximum of 80 kilos without snapping. Or Hugo is placed on a balance scale to weigh out 79 kilos of gold (for some weighty purpose). The question is whether Hugo ought to step on to the bridge, or whether that really is 79 kilos of gold, and not more. Odile volunteers, 'Hugo

²See Searle (1978, 1980) and Travis (1997, 2000, 2008).

weighs 79 kilos.’ Again, the claim is that Odile spoke truth in the first case and falsehood in the second. (Travis 1989, p. 19)

Suppose that the refrigerator is devoid of milk except for a puddle of milk at the bottom of it. Now consider two possible speakings, by Odile, of the words, ‘There’s milk in the refrigerator.’ For the first, Hugo is seated at the breakfast table, reading the paper, and from time to time looking dejectedly (but meaningfully) at his cup of black coffee, which he is idly stirring with a spoon. Odile volunteers, ‘There’s milk in the refrigerator.’ For the second, Hugo has been given the task of cleaning the refrigerator. He has just changed out of his house-cleaning garb, and is settling with satisfaction into his armchair, book and beverage in hand. Odile opens the refrigerator, looks in, closes it and sternly utters the above words. (Travis 1989, p. 18)

The argument goes as follows. If semantics determines truth-conditions, then the truth-conditions of the sentence should not vary. However, there are cases where truth-conditions do shift. To show this, Travis describes two scenarios differing in some assumption (what the conversation is about, the purposes of the interlocutors). So truth-conditions are not exhausted by semantics (plus syntax). As an alternative, Searle considers the possibility that some assumption we implicitly, in normal conditions, take to hold is explicitly rejected. For example, imagine that there is no gravitational force. In this case, it is not clear what the truth-conditions of a sentence as ‘The cat is on the mat’ would be. Which shows that the truth-conditions of the sentence are affected by background assumptions.

Travis’ examples are intended to show that semantics underdetermines truth-conditions. Truth-conditions depend on something beyond linguistic meaning. What is this something? Searle and Travis usually focus on different things, but we need to include here at least the following: assumptions about the world (like the assumptions about gravitational force), facts about how we do things, about our habits (we don’t drink coffee with stains of milk from the fridge), and specific features of the situation of use (what are we using the sentence for, what’s the goal of the conversation). As a consequence, understanding a language (grasping the satisfaction-conditions of an utterance) involves, besides semantic competence, implicit knowledge about the world, our society, and being aware of what is at stake in the situation of use.

Can we get rid of the underdeterminacy? A negative answer is supported by the fact that, although some of the elements that I have mentioned are easily encodable in a new, more complex sentence, this second sentence would in turn only make sense provided further assumptions are in place. Indeed, although Travis’ main argument is based on Travis cases, his radical view—namely, the view that natural language, and not only a few sentences, is underdetermined—is supported by the remark that the kind of examples he uses can be iterated, as he himself shows regarding the green leaves example. Suppose that one substitutes the sentences ‘The leaves are green’ by the more specific ‘The leaves are painted green’³:

³See Searle (1980) for similar remarks.

Or suppose [the leaves] are painted, but in pointillist style: from a decent distance they look green, but up close they look mottled. Is that a way of painting leaves green? It might sometimes, but only sometimes, so count. So there would be two distinct things to be said in the presumed 'paint counts' sense of 'is green'. (Travis 1997, p. 90).

The outcome of this more refined version of the argument is that one cannot simply get rid of the underdeterminacy by making explicit the variation between the two cases or somehow trying to select a sense, for further examples might be created for that new sense. What goes for the first sentence goes for any specification (in natural language) of the content this sentence expressed in a given occasion.

The conclusion that the number of examples offered motivates is that most (perhaps all) natural language sentences are underdetermined in the sense that their linguistic meanings do not fully determine the truth-conditions of their tokens. Travis uses the expression 'occasion-sensitivity' to refer to this form of dependence of truth-conditions (or extension in general) on particular occasions of use.

3 Two Versions of Truth-Conditional Pragmatics

Contextualism and occasionalism are versions of what is usually called truth-conditional pragmatics. According to proponents of truth-conditional pragmatics, truth-conditions cannot be recovered via syntactic and semantic elements alone. Rather, semantic content needs to be supplemented with contextual information. This approach contrasts with traditional truth-conditional semantic theories—theories that aim at attributing truth-conditions to sentences on the basis of their syntax and semantics.

Borg (2012) defines contextualism as the view that recovery of the truth-conditional content of an utterance at least sometimes calls for pragmatic interpretation. Contextualists claim that the semantics of at least some non-indexical sentences fail to determine a truth-conditional content. It is common to distinguish two kinds of contextualism, depending on the scope of the claim that semantics underdetermines truth-conditions. Moderate contextualism is the view that only some sentences exhibit this feature, whereas radical contextualism generalizes context-dependence. Here I am discussing Travis argument and its contextualist interpretation. Since examples similar to the ones that Travis himself has provided can be generated for a large class of sentence, acceptance of Travis cases leads to a radical form of contextualism, as Cappelen and Lepore (2005) argue. Radical contextualism as it is understood here is then similar to what Recanati (2004) calls the Wrong view format:

It's not just that there are expressions whose meaning is schematic and involves contextual variables; it is *linguistic meaning in general* which suffers from a form of underdeterminacy which makes it unfit to carry content save against a rich contextual background. Owing to that underdeterminacy, some form of enrichment or contextual elaboration becomes mandatory for the sentence to express a definite proposition. (Recanati 2004 p. 96).

Contextualism, as I understand it following Borg (2012) and as distinct from occasionalism, also has it that the content of an utterance is a proposition. As a result, a contextualist explanation of the phenomenon identified by Travis would involve two theses:

Underdeterminacy: semantics, or linguistic meaning, underdetermines truth-conditions.

Propositionalism: sentences, in contexts of use, express definite propositions.

Underdeterminacy, thus understood, concerns the relation between the linguistic meaning of a sentence and the truth-conditional contents it can be used to express. There are a variety of candidates to be ‘the proposition expressed’, so to speak, all compatible with the linguistic meaning of the sentence, and we need pragmatic interpretation in order to determine which one actually is, in an occasion of use, the content expressed by the sentence (the one the speaker intended, or the one determined by some contextual aspect, let’s say). If this content is thought of as a proposition, these two theses together imply an ineffability claim according to which sentences do not encode classical propositions. Propositions are, in this sense, ineffable.

Propositionalism distinguishes contextualism from occasionalism. Despite their agreement about the underdeterminacy of natural language, these two views are in disagreement when it comes to modelling the content of an utterance. Contextualism keeps the idea that there are propositions to be had; occasionalism has it that there are no representational items with intrinsic truth-conditions. The contextualist picture is endorsed by Carston:

It is fully propositional conceptual representations, rather than sentences, or even utterances of sentences, that are the primary bearers of truth conditions (Carston 2002, p. 60).

And is also suggested in the following passage:

‘Mine’, with its usual meaning, does not always refer to the property of being mine; why should ‘round’, with its usual meaning always refer to the property of being round? With this correction, Travis’ argument can be reconstructed as leading to the less radical conclusion that meaning alone ubiquitously underdetermines what is said. (Williamson 1998, p. 10 n, quoted in Travis 2008, p. 174)

However, the disagreement between contextualism and occasionalism only concerns some notions of propositions. Propositions have been conceived in two main ways: as structured items (strings of properties or concepts, let’s say) and as sets of possible worlds (truth-conditions). If the contextualist uses ‘proposition’ as a synonym for ‘truth-conditions’, then the distinction contextualism-occasionalism is blurred—an occasionalist as Travis takes utterances (sentences in context) to express truth-conditions. By contrast, if by ‘proposition’ he means a structured representational item with intrinsic truth-conditions (as something whose truth-conditions are independent of the tokening context, as a Fregean proposition), then there is a disagreement, for the occasionalist distrusts the assumption that strings of properties or concepts will behave differently from strings of words. Occasionalism takes the moral of Travis cases to be the claim that words have satisfaction conditions only when embedded in larger wholes, including a situation

of use and, in Searle's version, an indefinite number of background assumptions. Truth-conditions are relative to them.

Carston and Williamson's quotes also suggest two ways of fleshing out the idea that the content of our utterances is propositional in nature: one can model it as a proposition understood as a string of properties, or, one can model mental representations (strings of concepts) as being free from underdeterminacy (fully propositional) and take the content of an utterance to be the mental representation expressed by it.

In what follows I explore some reasons why one would not want to go for 'the less radical conclusion that meaning alone ubiquitously underdetermines what is said'—why one would rather avoid modelling the content of a predicate in an occasion of use as a property and that of a sentence as a string of properties or concepts.

4 Some Problems for Contextualism

4.1 *Travis-Borg Slippery Slope Argument*

Travis (1989, pp. 23–24) and Borg (2012, p. 36) argue against this approach. Travis notes that it can be objected to occasion-sensitivity that 'is green' or 'weighs 79 kilos' refer not to a property but to a family of properties. Let's suppose that this is the case and that, even if we cannot talk about the property that 'is green' or 'weighs 79 kilos' refer to, for there are a variety of them, we can talk about the property that the words 'is green' or 'weighs 79 kilos' refer to on a particular occasion, and assure a level of occasion-insensitive content. Travis argues that the assumption that our words-in-use, express an occasion-insensitive property is problematic because of the same reasons why the assumption that 'is green' expresses an occasion-insensitive property is problematic. In deciding that 'weighs 79 kilos' applies to Hugo in such and such conditions, or that 'There's milk in the fridge' is true of a certain fridge in such and such conditions, we have solved some doubts but not all of them. So suppose, for the sake of the argument, that there is a property that 'is green' and 'weighs 79 kilos' express in the scenario described by Travis. Call them P and Q. Now we can ask, would 'There's P in the fridge' be true if there was (only) a slice of cheese? Or synthetic milk? Or a bottle with coagulated milk? Would 'Hugo Qs' be true if the earth's gravitational force were to be halved overnight? We cannot answer these questions, for they raise doubts we didn't solve when we settled the issue about the bridge or about cleaning the fridge. So, Travis concludes, it turns out that these properties, in turn, leave some questions open—which, for him, means that they are occasion-sensitive after all.

Relying on Travis' argument, Borg claims that the proponent of the mixed view is under pressure to specify the occasion-insensitive contents that, according to him, are expressed by our utterances. The claim is that these contents are recovered via pragmatic interpretation, not that they cannot be had. However, it is reasonable to

think that for any specification of a given understanding that can be given, or for any sense we are able to somehow select, new Travis cases can be created and, as Borg says, ‘occasionalism at one remove is still occasionalism’ (2012, p. 17).

I think, however, that contextualists have the means to counter this argument, for they can reject the identification of the properties Travis specifies (what he calls P and Q) with the contents they postulate. As soon as we name or specify properties, we are dealing with linguistic meaning, and this admittedly underdetermines the property expressed. Thus, an advocate of contextualism might counter the Travis-Borg argument by carefully distinguishing metaphysical determination from what we can call the epistemology of content. He might have it that there is a property expressed by a use of ‘weighs 79 kilos’ but that normal speakers are unable to answer some questions about it. For instance, when we introduce a question the conversational context is updated, and this might affect our intuitions about the correct answer. When we are asked if a certain sentence would be true in certain conditions, it might be hard to tell whether our intuitions still track the proposition expressed in the first place, or rather the proposition the sentence in question would express in the scenario now described—whether, for example, we track the proposition that ‘There’s milk in the fridge’ would express as used by someone who keeps synthetic milk in his fridge. Hence, it can be argued that utterances express structured propositions that are free of underdeterminacy but that these are not fully transparent to us—either because we do not know everything about those properties or because when we reflect upon them we do it in new contexts, updated with new speech acts, and thus it is not warranted that we manage to refer exactly to the same property.

Nonetheless, I think that the Travis-Borg argument raises important points. It makes one wonder why the properties posited by the contextualist should escape underdeterminacy, given our inability to answer certain questions. In what follows I will argue that the contextualist assumption that there are propositions to be had gives rise to three mysteries. The first is why we cannot coin words for these underdeterminacy-free properties. Given our inability to coin words that encode fully determinate properties, the second mystery is why our mental representations should be any different from language. And if they are not different, then it remains to be explained how we manage to refer to ineffable properties. The third concerns the determination of the property expressed.

4.2 *Linguistic Mystery*

Contextualism implies a gap between what we can encode in natural language and what we can express with it. Since natural language suffers from semantic underdeterminacy, it seems that we cannot encode propositions. However, we can express them by using language—context is supposed to bridge the gap and allow us to infer what proposition a given utterance expresses. The first mystery concerns this gap between what we can encode and what we can express.

It is important to note is that the underdeterminacy of natural language cannot be overcome by coining new words. Let us go back to the example involving the word 'green'. On an occasion of use, the sentence 'The leaves are green' is used to describe certain leaves as, roughly, being superficially green. Suppose that we coin a new word, 'green_{sup}', that encodes this property, and add it to our vocabulary. There are reasons to think that the semantics of this word will in turn underdetermine the property expressed by it. As soon as we start using 'green_{sup}' in different contexts it will most likely behave like any other word in natural language and, in turn, underdetermine its satisfaction conditions on an occasion of use. Imagine that someone uses 'green_{sup}' to describe some other leaves. Now we need to decide how similar these new leaves have to be to Pia's leaves for the description 'The leaves are green_{sup}' to be true. But nothing prevents us from solving this question differently in different occasions. And if so, 'green_{sup}' will end up being as underdetermined as 'green'. More importantly, the kind of contextualism examined here involves a general underdeterminacy claim, so, unless somehow modified, it is part of its view that the new word will also exhibit semantic underdeterminacy.

Now, if uses of the word 'green' express occasion-insensitive properties, it should be possible to coin occasion-insensitive words, by which I mean words that are free of semantic underdeterminacy. The first unpalatable consequence that the contextualist needs to explain is why we cannot have words that encode (instead of simply expressing) occasion-insensitive properties.

4.3 *Cognitive Mystery*

The unpalatable consequence that I will introduce here has to do with the role that the ineffable properties are supposed to play in our cognitive lives. Here is a very simple picture that might sound plausible to those unconvinced by Travis cases. When Pia says 'The leaves are green' she intends to describe certain leaves as having the property GREEN—as being of a certain colour. She has a thought about the property GREEN; in order to convey it, she utters a sentence containing the word that encodes this property—'green'. Properties here can be used to model our cognitive lives. The support for the claim that Pia's thought is about the property GREEN can be provided by her linguistic behavior (she said 'green', so she must intend to describe the leaves as green).

The picture the contextualist is putting forward is a different one. According to contextualism, when Pia says 'The leaves are green' the word 'green' refers to a non-encodable property. But now, it remains to be explained how she manages to do so. How do these ineffable properties enter our cognitive lives? How do we refer to them? If conscious thought at least sometimes is linguistic, as introspection suggest, then there seems to be a gap between the property Pia could be consciously intending to attribute to the leaves (a certain colour), and the property that she in fact attributes. And this gap remains even if we describe Pia as thinking about an encodable but more fine-grained property, as PAINTED GREEN.

Nonetheless, a contextualist could offer additional reasons suggesting that our concepts correspond to the non-encodable properties posited. If so, it would be natural to model the content of our speech acts as structured strings of underdeterminacy-free properties. Carston, a prominent advocate of the view that natural language sentences underdetermine truth-conditional content and mental representations do not, argues that our concepts are unlike our words.⁴ Within her framework, concepts are not equivalent to natural language words—instead, they could be seen as referring to non-encodable, underdeterminacy-free, properties.

According to Carston, there is a gap between thought and language. Whereas our thoughts are completely propositional, our language does not have the means to encode them. The gap, however, is not necessarily difficult to bridge. In a context of use, a speaker will seek a sentence that, given the available contextual evidence, will allow the hearer to infer the thought he intends to communicate. So the hearer uses the sentence uttered, plus all the available clues, in order to grasp what the speaker intends to communicate.

Following Fodor, Carston understands having a thought as tokening a Mentalese sentence. The question is: is there some element that affects truth-conditions but that is not encoded in the Mentalese sentence, as happens with natural language? Carston's answer is negative, and relies on two elements of her theory. First, Carston assumes a computational view of mental states and mental processes, and this model involves a formality constraint: 'for every feature of content that a thought process is sensitive to, there must be a formal element present in the Mentalese representation of that content' (2002, p. 74). The underlying idea is that when it comes to inferential processing, the mind works with the Mentalese sentences it has. If, for instance (Carston's example), there is only one Mentalese concept corresponding to the word 'bank', then thoughts about financial institutions and about the side of the river would be indistinguishable. So Mentalese sentences must involve different concepts. This argument, if sound, applies to any form of equivocation, including the one present in Travis' examples. Second, public language is a medium for communicating thoughts, but one that 'evolved on the back, as it were, of an already well-developed cognitive capacity for forming hypotheses about the thoughts and intentions of others on the basis of their behaviour' (2002, p. 30). Natural language is an instrument for communicating thoughts, but is posterior to the ability to have thoughts and intentions and the pragmatic capacity to interpret others' thoughts. So why should we expect thoughts to inherit the underdeterminacy of natural language? In this model underdeterminacy is not a feature of representations, but of a (defective) instrument we have for sharing representations.

However, this position is not without problems. First, it is not clear that the Mentalese concepts Carston makes use of in her theory are that different from what we can encode in natural language. According to Carston's view, in the underdeterminacy scenarios (as Travis cases), interpreters construct ad hoc concepts:

⁴See especially her (2002).

The idea is that speakers can use a lexically encoded concept to communicate a distinct non-lexicalized (atomic) concept, which resembles the encoded one in that it shares elements of its logical and encyclopaedic entries, and that hearers can pragmatically infer the intended concept on the basis of the encoded one. The description of such concepts as ‘*ad hoc*’ reflects the fact that they are not linguistically given, but are constructed on-line (on the fly) in response to specific expectations of relevance raised in specific contexts (2002, p. 322).

As it is described, an *ad hoc* concept differs from the corresponding encoded concept in the specific information it includes. For Carston, concepts involve an encyclopaedic entry with information about the kind of object that the concept refers to—for example, in the case of the concept LEAF, that leaves are something that trees have, that they are green by the effect of chlorophyll, etc. An *ad hoc* concept LEAF* adds or eliminates some of this information—for example, about the colour green being caused by the effect of chlorophyll. Now, according to the underdeterminacy principle that Carston endorses, ‘the linguistic underdeterminacy of the proposition expressed by an utterance is an essential feature of natural language’. Carston supports this claim, among others, through Travis cases. If the linguistic meaning of a sentence underdetermines the proposition expressed by an utterance of it, then there is no such a thing as the proposition expressed by the sentence *qua* type. In the case of predicates, this amounts to claiming that there is no such a thing as the property that a predicate *qua* type refers to, for its linguistic meaning of underdetermines it. But now, if the only difference between encoded and *ad hoc* concepts is which features are included in the encyclopaedic entry, then there is no reason to think that *ad hoc* concepts will not suffer from underdeterminacy. In particular, it should not be expected that they are immune to further Travis cases.

A related problem arises if we consider Searle’s idea that satisfaction conditions are relative to a background. The question here is why should Mentalese sentences, or mental representations, be immune to this kind of relativity. In order to address this issue, it is important to be clear about what is the background supposed to include. In view of the relativity that Searle and Travis cases put forward, it seems that the satisfaction conditions of sentences are at least relative to assumptions about the world (general assumptions) and about the particular occasion of use (local assumptions). Among the latter I include, mainly, the goal or purposes of the conversation. I grant that *ad hoc* concepts might be so precise that the goal of the conversation no longer plays a role, because it somehow includes the relevant dimension (as BACHELOR* includes ‘eligible for marriage’ in a conversation about dating men, in one of Carston’s examples). However, this is not enough to eliminate underdeterminacy. These features only have satisfaction conditions relative to a lot of assumptions about the world—in this case, about a certain institution of marriage, according to which certain men can get married (note that who is eligible for marriage depends on the institution: a married man is still eligible for married in some countries).

Besides this, one might question whether it is necessary, as Carston assumes, that at the mental level there be no equivocation (as her argument for there being two concepts for ‘bank’ aims at establishing). Occasion-insensitive concepts are not the

only things that could resolve equivocation.⁵ Here is an example.⁶ Let's imagine a smooth conversation about red apples, with no misunderstandings. A is hungry, sees some apples but, due to the bad illumination conditions, is not sure whether they are red. She says 'I'm hungry. Are those red apples?' B answers 'No, they are green, and rotten. Don't eat them.' Should we assume that, when A wondered whether the apples were red and was trying to find it out she was tokening not the concept RED, together with the concepts APPLE, HUNGRY, and so on, but a concept created ad hoc to fit her situation, not equivalent to a composition of lexically encoded concepts? I think the hypothesis that we token ad hoc concepts lacks ground, and that there is an alternative explanation. When A arrives home and asks 'Are those red apples?' she forms a thought involving the concepts RED and APPLE (and perhaps an indexical concept corresponding to 'those'). Besides this, she activates a lot of information about the surroundings, including information about her state (she is hungry) and therefore about the goal of her question (to get food), information about what is edible, about her tastes, about apples, etc. All this information constrains the satisfaction conditions of RED APPLES to the effect that it applies only to a specific kind of apple. As a consequence, there is no reason to assume that concepts refer to the ineffable properties posited by contextualism.

4.4 *Metaphysical Mystery*

Another version of contextualism has it that, in Travis cases, two different properties are expressed. It remains to be explained what determines which property is expressed. When Pia utters 'The leaves are green', there are a variety of properties compatible with the linguistic meaning of 'is green': a property that only leaves painted on both sides have or a property that leaves have if they are green at least on one side; a property that only leaves painted with green paint have or a property that all looking green leaves have; a property that only leaves that look green on normal illumination conditions have or a property that all looking green objects have on precisely the time of utterance (some of which might look green because of the illumination conditions); the property of being green on, at least, 95% of the surface, or at least in 90%, and so on. In this section I will argue that neither the world nor the speaker's intentions can select among the potential candidates—unless one admits what I will call 'partial properties'.

First, the world. One could think that the world itself determines which property Pia's utterance is about. If, for instance, the relevant leaves are painted green, not naturally green, then that is the property at stake. Against this option, Belleri (2014) imagines a case where the relevant leaf is both naturally green and painted green. Here, the world fails to decide which of these two properties is the one

⁵See Recanati (2007) for a similar example.

⁶I adapt an example from Wieland (2010).

expressed. Pia's communicative intentions might decide instead. In this example, we can imagine Pia being interested only in naturally green leaves. However, there are other cases where intentions are not precise enough. When Pia, after painting the leaves, says 'The leaves are green', her intentions determine that painted green objects count as 'green'. But they need not decide exactly how the objects must be painted in order to count as 'green'—for example, how much of the surface must be green, or what style of painting has been used. Pia might very well have formed the desire to have her leaves are those of her neighbor, without having a very detailed idea of how the leaves should look like. If someone to ask her 'What about leaves painted as in pointillism? Is that good enough for you?', she could simply answer that she does not know. Linguistic meaning (a certain colour, in this case), plus some vague purposes and a sample of an object that satisfies the predicate in that occasion are insufficient to choose one property among a set of candidates.

Still, one could try to get rid of some of these candidates by having recourse to some form of dispositionalism. For instance, let's suppose that we have two competitors for being the property Pia referred to, P_1 and P_2 , and one object O that is P_1 and is not P_2 . Then, we could determine which property Pia actually referred to by making use of her dispositions: if, when Pia is shown O she judges that O is 'P' (where P is the generic word for both properties), then the property she expressed is P_1 . Otherwise, it is P_2 . The problem with this option is that Pia's judgements cannot be thus used, for any answer Pia could give to the question whether it is true that 'O is P' would be an answer about the content of 'O is P' in this new context, so not necessarily an answer about P_1 and P_2 , but perhaps about P_3 .

All this suggests that predicates in use refer only to what we could call partial properties. Partial properties are partial in the sense that they determine partial functions from objects to truth-values. For many objects, the function does not determine any truth-value.

4.5 Conclusions

I have presented some arguments against the contextualist approach. Although there is no knockdown argument against this kind of view, I think the preceding sections encourage those convinced by Travis cases and similar arguments to look for new approaches. Paraphrasing Wittgenstein, we can say that the contextualist has no model for his inordinate properties, but that he is seduced into using a super-expression.⁷

In the last subsection I considered the option that we replace a traditional notion of property—as something that establishes an exhaustive partition among objects of a certain domain—with a notion of property that allows for some partiality. I will

⁷'You have no model of this inordinate fact, but you are seduced into using a super-expression. (It might be called a philosophical superlative.)' Wittgenstein (2009, paragraph 192).

suggest now that this is the first step towards a radical pragmatist approach along the lines of Travis’.

Suppose that we model the content of Pia’s utterance as expressing a property that determines a partial function. For some leaves, the function does not yield a truth-value. For example, for leaves that have been painted with yellow and blue dots (such a leaf only looks green at a certain distance). Now we could sharpen this property in different ways. We could establish that leaves that look green at a certain distance count as having that property. Or that only leaves painted with green paint do. And so on. In this sense, one could claim, with Travis, that:

[T]here is in principle no end of opportunities for circumstances of a stating to matter to what was stated. There is no point at which circumstances choose for us some truth-evaluable item which is itself immune in principle to admitting of different further understandings. (Travis 2008, p. 6).

The context of use, in this case perhaps Pia’s intentions, determines what is relevant for whatever is at stake—the desire to have a beautiful garden, let’s say—but they need not extend to non-relevant cases—to other similar but not identical trees, or to leaves painted in other ways. For ordinary purposes, this might be good enough.

5 Occasion-Sensitive Languages

Let me finish by exploring what would occasionalism amount to. Occasionalism is an approach to language that challenges the once standard identification of meaning with extension. Travis cases invite the following line of reasoning: if linguistic meaning does not determine truth-conditions, how is it that utterances are truth-evaluable? There must be something else, in the context of use that, together with linguistic meaning, determines truth-conditions. In the contextualist approach, there seems to be a gap between conventional meaning and truth-conditions that needs to be bridged by some other content (a thought, a full proposition). The occasionalist approach, by contrast, does not have recourse to a further content. Rather, it takes it that sentences are truth-evaluable when embedded in the appropriate context—one that includes activities, purposes . . .

Occasionalism involves two claims:

Meaning: Linguistic meaning is merely preparatory. It is not to be identified with application conditions (or extension).

Application conditions: Words have application conditions in the context of broader activities.

Words as ‘green’ are used not only in the context of different sentences but also in the context of different activities. In these activities, they are understood differently. When not embedded in an activity, they do not have application conditions.

That words get a definite content when embedded in a context of use is revealed by two phenomena. First, we are unable to understand some well-formed sentences. Second, we have the ability to understand many sentences in a variety of ways, including new ways. Concerning our inability to interpret some well-formed sentences, Searle (1980) proposed the example 'Cut the sun'. We know what the expression types 'cut' and 'sun' mean. However, we are unable to interpret what are we supposed to do. Note that this has nothing to do with the fact that the sun is a huge object out of reach: we could imagine a giant cake and, although nobody would be in a position to satisfy the order 'Cut the cake', we know what it would demand of us. What is interesting is that, if we provide enough contextual information, we arrive to an interpretation. For instance, in a science fiction novel, a procedure to cut stars into two could be described. Readers would then interpret the order 'Cut the sun' as calling for the application of this procedure. This shows that what was missing was a proper context in which to embed the utterance. If we create the appropriate context (the activity in which to use the words), then we might be able to understand the sentence.

Second, if we imagine a new activity in which to use a sentence, we can create a new way to understand it. Thus, for instance, even if as a rule we use colour predicates to describe the way things look in a variety of conditions, we can easily imagine situations in which colour predicates describe how things will look, not how they look in the moment of utterance. For instance, imagine some people whose work consists in dyeing clothes. This people might describe some piles of white cotton clothing not as 'white' but as 'blue' if that is the colour the clothing from that pile will have. One of them might ask 'Where is the blue clothing? I want to check the quality of the cotton.' The answer 'Those are the blue ones' will be true if those are the ones that will be dyed blue.

I will finish with some remarks about communication. In the last years, the contextualist debate has shown that, even if semantics delivered a truth-evaluable proposition, these so-called minimal propositions cannot ground communication. They are insufficient. The reason is that what is relevant for communication is the contextually modulated content, not the minimal proposition. Going back to Pia's example, successful communication would require that both Pia and the botanist understand 'green' in the same way—as equivalent to 'naturally green', let's say. Otherwise, they would misinterpret each other's utterances.

Nothing in this example forces us to conclude that we need more precise propositions in order to secure communication. What we need is to assure convergence on interpretation. Transmission of an ineffable proposition might be one way to achieve it. But we can explain convergence without recourse to ineffable propositions. There is convergence when Pia and the botanist take the same kind of leaves (for example, those in Pia's tree) to satisfy, or not to satisfy, the predicate 'is green'. My suggestion is that communication requires that, given a situation or occasion, interlocutors assign the same satisfaction conditions to utterances: that they see the same things as fitting a given description and the same actions as being permitted or forbid by a given utterance (as being in accordance with an assertion, as being a fulfilment of a promise, and so on). This view bears some similarities with an approach put

forward by Davidson. Davidson (1986) distinguishes prior from passing theories of language. A theory of language is a recursive theory that assigns truth-conditions to well-formed expressions of a language. Prior theories are the truth-conditional theories interlocutors are equipped with when they face a new conversation. Passing theories are the theories used in interpreting the utterances of the conversation, and they need not be equivalent to prior theories. Communication succeeds when the interpreters' (speaker and hearer's) passing theories converge. We can see passing theories as theories that are adjusted to the context of use. In a conversation about science, the interlocutors' passing theories will restrict the satisfaction conditions of 'is green' to things that are naturally green.

Convergence on attribution of satisfaction conditions in an occasion of use is warranted by two elements. First, words have a certain linguistic meaning. This linguistic meaning imposes some constraints. For instance, the word 'green' refers to a particular colour (not red, not blue, etc.) Second, speakers of a given language share both an indefinite number of assumptions about the world, human and cultural habits, abilities and so on and a common sensibility, so to speak, by which I mean the ability to see things in similar ways. Being able to learn a language involves sharing certain psychological traits with the other speakers. As Wittgenstein argued, any ostensive sign, as well as any series of examples, is compatible with indefinitely many interpretations. We, however, converge with the rest of speaker on an interpretation and thus are able to share a language. Thus, when a group of interlocutors engage in a conversation, they already share a common sensibility: this is guaranteed by the fact that they all speak the same language.⁸

Here is a rough account of what goes on in a conversation. Conversations happen at particular situations involving objects, places, times . . . Moreover, conversations involve broader, not necessarily linguistic activities, as discussing about a particular topic, telling jokes, shopping, trying to find out whether there's food around . . . Interlocutors are, so to speak, situated in these situations and activities. Very often they agree on what is the aim of the conversation, its topic. Suppose, going back to a previous example, that we are talking about having a coffee. Assuming some common knowledge about coffee habits, we would probably take some milk stains in the fridge not to satisfy 'milk'. But for that, we do not need an ad hoc concept MILK*: we just need to be aware of what is going on and infer from all we know that 'milk' here applies to milk suitable for coffee. The passing theory used here is one in which stains in a fridge do not make the sentence 'There's milk in the fridge' true.⁹

⁸See Recanati (2010) for a similar argument.

⁹This work has been funded by the research project "Objectivity-Subjectivity in Knowledge and Singular Representation", FFI2015-63892-P (MINECO/FEDER, UE).

References

- Belleri, D. (2014). You can say what you think: Vindicating the effability of our thoughts. *Synthese*, 191(18), 4431–4450.
- Borg, E. (2012). *Pursuing meaning*. Oxford: Oxford University Press.
- Cappelen, H., & Lepore, E. (2005). *Insensitive semantics: A defense of semantic minimalism and speech act pluralism*. Oxford: Wiley-Blackwell.
- Carston, R. (2002). *Thoughts and utterances*. Oxford: Blackwell.
- Davidson, D. (1986). A nice derangement of epitaphs. In E. Lepore (Ed.), *Truth and interpretation: Perspectives on the philosophy of Donald Davidson* (pp. 433–446). Oxford: Blackwell.
- Recanati, F. (2004). *Literal meaning*. Cambridge: Cambridge University Press.
- Recanati, F. (2007). *Perspectival thought: A Plea for moderate relativism*. Oxford: Oxford University Press.
- Recanati, F. (2010). *Truth-conditional pragmatics*. Oxford: Oxford University Press.
- Searle, J. (1978). Literal meaning. *Erkenntnis*, 13(1), 207–224.
- Searle, J. (1980). The background of meaning. In J. R. Searle, F. Kiefer, & M. Bierwisch (Eds.), *Speech act theory and pragmatics* (pp. 221–232). Dordrecht: Springer.
- Travis, C. (1989). *The uses of sense: Wittgenstein's philosophy of language*. Oxford: Oxford University Press.
- Travis, C. (1997). Pragmatics. In B. Hale & C. Wright (Eds.), *A companion to the philosophy of language* (pp. 87–107). Oxford: Blackwell.
- Travis, C. (2000). *Unshadowed thought: Representation in thought and language*. Cambridge, MA: Harvard University Press.
- Travis, C. (2008). *Occasion-sensitivity: Selected essays*. Oxford: Oxford University Press.
- Wieland, N. (2010). Context sensitivity and indirect reports. *Philosophy and Phenomenological Research*, 81(1), 40–48.
- Williamson, T. (1998). Indefinite extensibility. *Grazer Philosophische Studien*, 55, 1–24.
- Wittgenstein, L. (2009). *Philosophical Investigations* (4th edn) (Hacker and Schulte, Trans.). Oxford: Wiley-Blackwell.

Context and Communicative Success



Joey Pollock

Abstract Traditional accounts of the conditions on communicative success are invariantist. For example, some authors claim that, for communication to succeed, a hearer must always grasp the very content that the speaker expressed with her utterance; others claim that success is always proportional to the degree to which the hearer understands this content. In this paper, I argue that these invariantist approaches cannot offer a comprehensive account of communicative success. When we attempt to communicate, it is usually with the intention to secure some perlocutionary effect beyond the communicative exchange itself (for example, to get our interlocutor to pass the salt, or to convince her to form a particular belief). I argue that, relative to these intentions, it may not matter whether a hearer perfectly understands the speaker; it may not even matter if the hearer's understanding is quite poor overall. In place of invariantist conditions (or perhaps, in addition), I propose an approach that is context-dependent in the following sense: holding fixed the content expressed by the speaker, differences in features of the broad context of the speech exchange can determine differences in the standards for communicative success. On this approach, success requires that the hearer understand the content expressed by the speaker in ways that are relevant to the speaker's perlocutionary intentions.

Keywords Communication · Context · Perlocutionary intention · Understanding · Mental content

J. Pollock (✉)
ConceptLab, University of Oslo, Oslo, Norway

© Springer Nature Switzerland AG 2020
T. Ciecierski, P. Grabarczyk (eds.), *The Architecture of Context and Context-Sensitivity*, Studies in Linguistics and Philosophy 103,
https://doi.org/10.1007/978-3-030-34485-6_13

245

1 Introduction

When a speaker produces an utterance, under what conditions does she succeed in communicating with her audience?¹ There are different answers to this question in the literature. For example, some authors adopt a ‘shared content’ approach to communication. Shared content approaches claim that, for communication to succeed, the hearer must grasp the very content that the speaker expressed with her utterance. An alternative approach is to claim that success requires, not shared content, but merely some suitable degree of similarity between the content grasped by the hearer and that expressed by the speaker. Whilst these approaches offer different pictures of communicative success, they are each committed to a kind of invariantism: they claim that the standards for communicative success remain constant across different speech exchanges. In this paper, my aim is to demonstrate that there are instances of successful communication that these invariantist approaches cannot accommodate. What is missing from invariantism is an appreciation of the ways in which features of the context of the communicative exchange can place varying demands on interlocutors. I will argue that consideration of the role of such features in communication motivates a context-dependent understanding of communicative success (either alongside, or instead of, invariantist conditions). It is common to appeal to context in order to explain various elements of a communicative event. For example, relevance theorists have argued that context plays an important role in the processes by which a hearer recovers content from the speech exchange; context is also thought to play some role in determining what is said by the speaker for some kinds of expression. My approach in this paper highlights a different kind of function that certain features of context (broadly construed) can play in a communicative exchange: that of determining the standards for communicative success.

The paper proceeds as follows. In Sect. 2, I offer some preliminary remarks concerning communicative exchanges and content individuation. In Sect. 3, I explain several invariantist conditions on successful communication. In Sect. 4, I offer a series of examples designed to motivate my context-dependent account over (pure) invariantism. In Sect. 5, I consider further support for the intuitive verdicts on these examples.

2 Communication and Content Individuation

A communicative exchange is an event with the following structure: the event begins with some contentful mental state of the speaker; this mental state (partly) causes the production of an utterance; the utterance is received by a hearer; and this hearer

¹I understand ‘utterance’ broadly here, to refer to the signal in a communicative exchange; I intend this to include signals that are written, signed, spoken aloud, etc. Similarly, I use ‘speaker’ and ‘hearer’ to refer to the sender and receiver of this signal, respectively. (cf. Grice 1957: 380)

is thereby caused to enter into a contentful mental state of her own.² Beyond this rather simple description of the event, there are many questions that arise concerning the finer details. For example, the above description involves two mental states of the speaker and hearer, respectively; how is the content of these states individuated? What relation does this mental content stand in to the content of the utterance that the speaker produces? How is it exactly that the hearer recovers content from the exchange? What we say about the nature of mental and utterance content can make a significant difference to an explanation of communication and communicative success. Thus, before examining various candidate conditions on communicative success, I will first say a bit about content individuation.

It will be useful to make a distinction between sociolectal and idiolectal theories of content individuation (for both mental content and utterance content³). Sociolectal theories claim that subjects can very often think thoughts, and produce utterances, that contain or express concepts that are type-identical to those possessed by others in the same language community. That is (ambiguity and polysemy aside), speakers in a language community who utter the same word-forms will often thereby express the same concepts. Many externalist approaches to meaning and thought content are sociolectal (e.g., Burge 1979; Putnam 1975; Fodor 1994; Millikan 1984). Externalists claim that content is individuated, in part, by factors external to the subject, such as her social and/or physical environment. In contrast, idiolectal theories claim that, rather than sharing a public language, subjects each speak and think in their own idiosyncratic idiolects. On idiolectal theories, two subjects will rarely express the same concepts with utterances of the same word-forms. This approach is most often internalist; internalists claim that content is determined solely by factors internal to the individual. An example would be a radical holism according to which concepts are individuated by their relationships to all other concepts in a subject's conceptual web (e.g., Rapaport 2003).

One central difference between the approaches is how they treat the relationship between content and (a certain kind of) conceptual and linguistic understanding. A sociolectal theory will allow that content and understanding can come apart: a subject can grasp a concept whilst having very little understanding of its content. To take Burge's (1979) famous example: a patient, Alf, is said to possess the communal concept ARTHRITIS even though he is seriously confused about the concept's correct application. Alf believes that he has arthritis (a disease of the joints only) in his thigh; and this, Burge claims, is a conceptual error. Subjects like Alf, who are conceptually confused, can nonetheless share concepts with others in their community because the conditions on conceptual grasp are relatively undemanding. Burge suggests that, to be attributed a communal concept, a subject must possess

²Here I am roughly following Pagin (2008). Pagin offers a more detailed, and more general, characterisation of the structure of a communicative event; however, many of these details are not important for my purposes.

³These theories of content individuation are also relevant to the determination of the linguistic meaning of words and sentences, where linguistic meaning can be thought of as distinct from utterance content.

a general competence with respect to her language as a whole (or relevant parts of it) and be disposed to defer to experts with respect to the correct application of the concept (1979: 114). Other than this, however, subjects can be quite seriously mistaken – or just agnostic – as to the application conditions of a concept and still possess it. On sociolectical theories, once we have settled the question of whether an individual grasps a particular concept, we have not thereby settled the question of whether (or to what extent) she understands this concept. As I will explain below, this means that one might want to require *both* content recovery *and* understanding in one’s account of communicative success. Whereas a sociolectical theory will allow that understanding can float largely free of content, an idiolectical theory will claim that content closely tracks our understanding of it: we cannot be very wrong about the content of the concepts we possess. For example, if I am disposed to assent to ‘All dogs are molluscs’, an idiolectical theory might claim that I don’t have incorrect understanding of a public DOG concept; rather, I have *correct* understanding of my own idiosyncratic concept. On this approach, content and understanding are just two sides of the same coin. For an idiolectical theory, then, grasping a content, and understanding that content, might come to more or less the same thing.

The structure of the account of communicative success that I will argue for in this paper is designed to be compatible with many different theories of the individuation and structure of content, including both sociolectical and idiolectical theories; however, the way that the account is stated will vary significantly depending on one’s preferred approach to these issues. For ease of exposition, I will make a few assumptions. Firstly, I will assume the sociolectical framework for both thought content and utterance content when presenting my argument. Secondly, I will restrict my focus to exchanges in which an utterance’s content does not diverge from the thought content it is used to express.⁴ Thirdly, I will assume that an utterance can express, at most, *one* content – that is, I assume what Cappelen and Lepore (2005) call ‘speech act monism’. In the next section, I explain in more detail the invariantist conditions on communicative success that I will critique.

3 Communicative Success and Invariantism

Invariantist theories of communicative success claim that the standards for success do not vary with the context of the speech exchange – where ‘context’ here, should be understood in the broad sense, to include things like the interlocutors’

⁴Utterance content need not always match the content of the thought expressed – subjects may sometimes misspeak, for example. Cases in which the two diverge cause complications. I think that grasp of utterance content itself is not required for communicative success. In Pollock (2015), I argue that, on sociolectical theories, what must be grasped by a hearer is not utterance content, nor even thought content, but a speaker’s *understanding* of the thought content she expresses (Thanks also to Andrew Peet for helpful discussion of this issue). I ignore this complication in what follows.

background knowledge⁵ as well as physical features of the environment of the speech exchange. For the invariantist, this context might play some role in, for example, the hearer's *recovery* of content; but it plays no role in determining the standards for communicative success in the following sense: if a speaker expresses the very same thought content, on two different occasions, the standards for success will remain constant across both exchanges. In contrast, I will understand a context-dependent approach as one which claims that, holding fixed the content expressed by the speaker, differences in features of the broad context of the speech exchange can result in differences in the standards for success.⁶ I don't claim that this is the only way to understand what context-dependence amounts to with regards to communicative success; rather, this understanding marks an interesting division between different kinds of approach.

There is a range of invariantist conditions on communicative success that you might endorse. In the next section, I will set out two prominent kinds of necessary condition: conditions on content, and conditions on understanding.⁷

3.1 *Conditions on Content*

When setting out the structure of a communicative event, I said that the exchange includes two contentful mental states of the speaker and hearer, respectively. One kind of condition on communicative success claims that communication requires that some special semantic relationship hold between these mental states. As mentioned above, there are different options for what this relationship must be. The 'same content' approach claims that the content recovered by the hearer must be type-identical to the thought content expressed by the speaker. This is a relatively popular approach in the literature. It is most common when focus is restricted to testimonial knowledge acquisition (E.g., Burge 1993; Fricker 2006; and Goldberg 2007). The approach is also somewhat popular for communication in general; although it is more controversial whether *all* kinds of communication require shared content. Goldberg (2007), for example, allows that non-testimonial communication might not be so demanding. However, there are those who think that sharing content is central to understanding even non-testimonial communicative success. Cappelen and Lepore (2007: 115), for example, write that "Content sharing [. . .] is the very foundation of communication." (another example is Newman 2005). If you think

⁵We could think of this as common ground (Stalnaker 2002; Clark 1996). However, I don't think it is necessary that speaker and hearer share the relevant beliefs.

⁶Note that the account is not contextualism about the *term* 'communicative success'.

⁷These aren't the only conditions you might include. For example, perhaps the hearer must recover whatever attitude the speaker takes towards the content she expresses, as well as things like the illocutionary force of the utterance.

that the shared content approach is too demanding, there are various similarity-based alternatives in the literature. For example, one could say that communication succeeds only if the hearer grasps a content that is highly similar to the thought content expressed by the speaker; or that communication succeeds to the degree that the content grasped by the hearer is similar to the content expressed by the speaker. One example of a defender of content similarity is Bezuidenhout (1997) (other examples are Pollock 2015; Jorgensen 2009; Rapaport 2003; Carston 2001). She writes,

We need recognize only speaker-relative utterance content and listener-relative utterance content and a relation of similarity holding between these two contents (where the speaker-relative utterance content is identical to the content of the thought the speaker expresses and the listener-relative utterance content is identical to the content of the thought the listener comes to entertain). (Bezuidenhout 1997: 212)

These various content conditions on success are all invariantist. For example, shared content approaches claim that a hearer must always grasp the very content that the speaker expressed; graded theories claim that success is always proportional to the degree of similarity that holds between the contents expressed by the speaker and grasped by the hearer. It might be that different kinds of communication, or different types of utterance, involve different kinds of condition on communicative success; however, to claim this is not to claim that the standards for success vary with context.

3.2 *Conditions on Understanding*

In addition to these content conditions, it seems plausible that successful communication requires some kind of linguistic or conceptual understanding on the part of the hearer. There are many theories in the literature concerning what linguistic understanding consists in.⁸ Only some varieties of understanding would be sensible candidates for *additional* conditions on communicative success. For example, according to one approach, an act of linguistic understanding is just the means by which content is recovered from the speech exchange. Goldberg (2007), for example, proposes a notion of understanding that characterises it as the process by which a hearer maps the speaker's utterance onto an item of mental content in her own cognitive economy. On this approach, to achieve a state of correct understanding *just is* to have grasped the content that the speaker expressed.⁹

⁸Longworth (2010) examines several options.

⁹This approach could be framed as a traditional propositional theory of linguistic understanding, according to which understanding an utterance, *u*, amounts to knowing that *u* means that *p*, where *p* is some correct specification of *u*'s content (e.g., Heck 1995; Evans 1982; Dummett 1978); however, I think it is also compatible with weaker theories that do not treat understanding in terms of propositional knowledge about what was said (e.g., Longworth 2018).

Following Goldberg, I will call this kind of understanding ‘comprehension’ to distinguish it from other varieties of linguistic understanding.¹⁰ If you hold this kind of theory (and you think there is no other kind of understanding) then you might have no need for both a condition on content and a condition on understanding in your account of communicative success.

When might one want a separate understanding requirement? Recall that, in Sect. 2, I explained that sociolectical theories of content entail that subjects can grasp contents that they do not understand. Comprehension, as described above, is the kind of understanding involved in grasping, or recovering, a particular content from a speech exchange. As already emphasised, this still leaves open whether (or to what extent) the hearer understands the content that she has recovered, in some further sense of the term. This latter kind of understanding, I will call ‘inferential’ understanding.¹¹ On this approach, understanding utterance or thought content is a matter of grasping the inferential relations that hold between this content and the rest of your beliefs. You understand the content of an utterance or thought correctly *to the degree* that you grasp the inferences that it is correct to draw from it. Importantly, inferential understanding is multidimensional: because contents stand in a great many inferential relations to other contents, there are different ways in which you can understand the same content that fall short of a full, correct, understanding. You could, for example, understand the content in a way that is just plain wrong (e.g., if you are disposed to infer from ‘Some dogs bark’ to ‘Some molluscs bark’). Often, however, what happens is that we understand contents partially. For example, you might be disposed to infer from ‘Stanley is an octopus’ to ‘Stanley lives in the sea’; but perhaps not from ‘Stanley is an octopus’ to ‘Stanley has 3 hearts’. Something like this notion of understanding must be in play in sociolectical theories in order to explain the sense, highlighted above, in which a subject can misunderstand the content of her own thoughts and the utterances of others, despite succeeding in grasping these contents. Later, I will try to convince you that this notion of understanding is also integral to an account of communicative success.

An inferential understanding requirement can be fleshed out in different ways, which more or less mirror the content conditions introduced above: success might require that a hearer fully understand the content that she recovers from the exchange, or that she understand quite well; or perhaps she succeeds to the degree that she understands the content recovered. Like the content conditions, these understanding requirements are invariantist and for the same reasons. It might be that no one of the above conditions (on content and on understanding) alone would be sufficient for communicative success. You can hold many of them together if you like. For example, you could claim that the hearer must grasp (comprehend)

¹⁰Goldberg uses the phrase ‘reliable comprehension’; this is because he is interested in how hearers can acquire testimonial knowledge. As my focus here is not testimony, the reliability dimension of Goldberg’s account is not important for present purposes.

¹¹Inferential understanding should not to be confused with the inferential model of communication embraced by relevance theorists. The latter is a view about the role of inference in arriving at an interpretation of the speaker’s utterance (Sperber and Wilson 1986; Bezuidenhout 1998).

the right content and also achieve full (inferential) understanding of this content. This is a pretty demanding approach. We could strengthen it still further: perhaps the hearer must not only grasp the right content but *know* that she has done so and *know* that her (inferential) understanding is correct. I am inclined to think that such a requirement is too strong – successful communication can be, and often is, lucky (cf. Davidson 1986)¹²; however, the account that I go on to present is compatible with these strengthened requirements on communicative success.

So, to summarise where we've got to so far: A communicative attempt is an event in which a speaker produces an utterance that expresses a thought content; this utterance is received by a hearer who, as a result, enters into her own contentful mental state. For ease of exposition, I'm assuming that some sociolectical theory is true of both thought and utterance content; I am restricting my focus to cases in which these two kinds of content do not diverge; and I'm assuming that an utterance expresses, at most, one content. I've presented some invariantist conditions on communicative success; the important ones are conditions on content and conditions on inferential understanding. On sociolectical theories, when a hearer observes the speaker's utterance, she both *grasps* a content, and (inferentially) understands this content in a particular way. I will now argue that the invariantist conditions – no matter their combination – cannot offer a plausible (or, at least, a comprehensive) picture of communicative success.

4 Against Invariantism

When we communicate, we are often trying to get our audience to do something – either for us, or with us. For example, to pass the salt, to lift one end of a piano, to form a particular belief, etc. I suspect that, on a suitably broad understanding of what it is to get people to do things, this is the only aim of communication. However, I will not argue for this claim here. Instead, my aim in this paper is somewhat modest: I claim that there is a kind of communication, the aim of which is to get people to do things; and the standards for success, for this kind of communication, are context-dependent. As such, I do not claim that invariantist conditions are never relevant to communicative success. There may be some cases, or kinds, of communication for which all success requires is that you grasp the right content, for example – if so, we should be pluralists about communicative success. My thesis here is that, when we consider the sorts of aims and intentions that a speaker can have, we will find that invariantist accounts cannot tell the whole story.

What features of context affect the standards for communicative success? In what follows, I will focus on what Austin (1962) calls the 'perlocutionary' intentions of the speaker. Perlocutionary intentions are intentions to produce perlocutionary effects. These are effects that utterances can have on the hearer which may fail to

¹²Pagin (2008) discusses several knowledge requirements on communicative success.

obtain even when ‘uptake’ is secured.¹³ Examples of perlocutionary effects include things like persuading, surprising, and amusing. Austin contrasts perlocutionary intentions and effects with ‘illocutionary’ intentions and effects. Illocutionary effects are those which are guaranteed so long as uptake is secured (and perhaps even without uptake).¹⁴ Paradigmatic examples are things like warning, promising, and requesting.

Illocutionary intentions are important for understanding *what it is* that the speaker intended to express with her utterance: they are intentions to produce an utterance with a particular content and force. Perlocutionary intentions, on the other hand, are important for understanding *why* a speaker attempts to communicate in the first place. Consider, for example, that a speaker may have the illocutionary intention to request that you ϕ ; but it is unlikely that this is her sole reason for attempting to communicate – nor would it provide a very illuminating explanation for her utterance. If I ask you to pass me the salt, I don’t ask you to do this (merely) because I want to ask you to pass me the salt; rather, I ask you this because *I want you to pass me the salt*. My intention to ask you to pass me the salt is my illocutionary intention; the intended perlocutionary effect, however, is that you actually pass me the salt. This picture of communication and its role in securing effects beyond the communicative exchange itself is broadly similar to Herbert Clark’s (1996) account, according to which communication is itself a form of joint action that is performed in service of further joint projects.^{15,16} Clark writes, for example:

Signals aren’t important merely because they mean things. They are important because they are used in discourse to accomplish the participants’ goals. When the server in the drugstore said “I’ll be right there,” she meant that she would be ready to serve me soon. But she was using the signal to coordinate her and my actions at that point in our transaction. Viewed in isolation, a signal is an act by which a speaker means something. Viewed within joint activities, it is an act by which the participants coordinate the next step in their ongoing activity. Signals are coordination devices. (1996: 132)

I doubt that many would deny that perlocutionary intentions are relevant to explaining why speakers attempt to communicate; however, existing accounts of communicative *success* often overlook perlocutionary intentions of the kind I am

¹³‘Uptake’ is Austin’s terminology. Uptake involves understanding the force and the content of an utterance.

¹⁴It is a controversial issue whether uptake is required for the success of an illocutionary act (E.g., Bird 2002; Langton 1993).

¹⁵Some aspects of my account might be in tension with Clark’s picture of communication. For example, I claim that the speaker’s perlocutionary intentions are central to determining the standards for communicative success; Clark (1996: 212), in contrast, emphasises how the aims of a communicative exchange are jointly developed and amended by the speaker and hearer throughout the course of a conversation. I think a version of my account could be made compatible with Clark’s approach, but it is beyond the scope of this paper to explore this issue.

¹⁶Tomasello (2008) also provides a thorough development of the conception of communication as central to the coordination of cooperative projects.

interested in here.¹⁷ I will argue that this kind of intention is important to determining the standards for communicative success. It is, in part, because these perlocutionary intentions vary from context to context that the standards for success are context-dependent.

What are the standards for success that change with context? I will argue that it is the manner in which the hearer must inferentially understand the content expressed by the speaker that is affected: different perlocutionary intentions can make different demands on the hearer with respect to how she must understand the content that the speaker expressed.¹⁸ Note here that, although perlocutionary intentions often aim at the hearer performing some action, my claim is *not* that success requires the performing of an action (such as passing the salt, or forming a belief). Success is not a matter of satisfying the speaker's perlocutionary intention (of producing the desired effect); rather, success is a matter of the hearer *understanding* the content recovered in a way that is determined by the speaker's perlocutionary intention: the kind of understanding required is the understanding necessary for the hearer to deliberately produce the intended effect as a result of the exchange (where she might achieve such understanding without then going on to produce this effect for various reasons). As such, my position is a kind of understanding condition on communicative success: it is a context-dependent version of the understanding requirements that were introduced in Sect. 3.2.¹⁹ How this is supposed to work will become clearer through consideration of examples.

In what follows, I will present three pairs of examples. In each pair, the first example will be an instance of communication failure and the second will be an instance of communicative success. At this stage, my claim is that these are the verdicts that intuition supports. I will have a bit more to say in justification of these intuitions in Sect. 5. Across each pair, the speaker will express the same thought content and produce an utterance with the same content; the hearer will grasp the same content each time. What *does* change across each pair are features of the context (broadly construed): there are changes in the projects that the interlocutors are engaged in. This includes changes to the perlocutionary intentions of the speaker and also the interlocutor's background beliefs. My thesis is that the differences in communicative success should be explained by appeal to these differences in

¹⁷When authors do consider perlocutionary intentions in relation to communicative success, it is usually to discuss whether communicative intentions include perlocutionary intentions of the sort present in Grice's (1957) account. This issue is discussed in, e.g., Recanati (1986) and Bach (1987). The perlocutionary intentions that I appeal to in my account are not communicative intentions in this sense; rather, they are intentions that aim at effects beyond the exchange itself. Moreover, I do not claim that communicative success requires their satisfaction.

¹⁸In framing things this way, I do not mean to suggest that communication failure is always or usually solely attributable to a hearer's failure to live up to these demands – the speaker is also responsible for choosing demands that are reasonable.

¹⁹I think the correct version of the theory is a graded version according to which communication succeeds to the degree that the hearer achieves understanding along the dimensions determined by the speaker's perlocutionary intentions. I will ignore this complication in what follows.

context: these features of context place different demands on the hearer with respect to the manner in which she must understand the content expressed by the speaker.

Example 1a

Sally and Hannah are at a party. Sally has noticed that Bill, who claims to be vegetarian, is eating shellfish. She utters (1) to Hannah with the perlocutionary intention that Hannah form the belief that Bill is eating shellfish.

(1) [Gesturing] He is eating shellfish.

Hannah knows what shellfish are; however, her view is obstructed and so she cannot tell who Sally is gesturing at.

In this example, Hannah is unable to grasp the content that Sally expressed with her utterance (although she does recover part of it); nor can she fully understand what Sally has said to her. An invariantist will treat this as a straightforward case in which communication fails because the hearer cannot recover the object referred to by the pronoun ‘He’.²⁰ And, indeed, it looks initially plausible to think that the moral to draw here is simply the unsurprising one that some kinds of communication require preservation of reference.²¹ However, I think that this need not be the case. Consider the next example.

Example 1b

Sally and Hannah are at a party. Sally knows that Hannah is severely allergic to shellfish. She utters (1) to Hannah with the intention that Hannah take steps to avoid getting sick (e.g., by leaving the room).

(1) [Gesturing] He is eating shellfish.

Hannah knows what shellfish are; however, her view is obstructed and so she cannot tell who Sally is gesturing at.

Once again, Hannah is unable to fully recover the content that Sally expressed. However, in this case, I think we should say that communication can succeed at least *relative to Sally’s perlocutionary intention*. Just like in (1a), nothing enables Hannah to identify the correct referent; however, in this example it doesn’t matter. Because Sally’s intention is different in this case, Hannah doesn’t need to know who she is gesturing at in order for communication to succeed. All that Hannah needs to understand is that someone is eating shellfish (or perhaps even just that there is shellfish present) – this much will allow her to recognise that she is in harm’s way – and she can understand this without grasping the referent of ‘He’ in Sally’s utterance. This suggests that, for at least some communicative purposes, it isn’t always required that we get the reference right. Communicative success in example (1a) does require identification of the referent; but success in (1b) is not so

²⁰For the ‘graded’ conditions on communicative success, this would be better described as a case in which communication is poor. The same thing goes for the other example pairs.

²¹There is a debate concerning the conditions on successful *de re* communication (e.g., Bezuidenhout 1997; Recanati 1993; Buchanan 2014). However, much of this debate focuses on the question of whether utterance content includes modes of presentation of the referents of referring expressions. As my focus here is on inferential understanding of thought content, my argument is somewhat orthogonal to this debate.

demanding. The degree of understanding required appears to be constrained by the speaker's perlocutionary intention in each case.²²

Example 2a

Stanley and Hank are at the zoo to take pictures for Hank's school project on flightless birds. Approaching a pen, Stanley utters (2):

(2) There is a cassowary in the pen.

Stanley utters (2) with the perlocutionary intention that Hank believe that there is a flightless bird in the pen that he might include in his project. However, although Hank understands most of the utterance, he does not fully understand the term 'cassowary'. He knows only that a cassowary is a dangerous animal.

Recall that we are assuming in these examples that Stanley and Hank share a sociolect. As such, it is reasonable to stipulate that Hank grasps the correct content.²³ We can also stipulate that Hank does not work out from the context of the exchange that (Stanley believes that) a cassowary is a flightless bird (although hearers often do repair exchanges in this way). Rather, this is an example in which Hank grasps the right content but does not understand it fully. In this example, I think we should say that communication fails (or is poor), again, at least relative to Stanley's perlocutionary intention. At most, Hank's understanding would allow him to infer that there is a dangerous animal in the pen – and this is not what Stanley intended that he do. As with example (1a), it might initially look like the problem here is simply that Hank doesn't understand Stanley's utterance well enough – the invariantist might say that his understanding of 'cassowary' is not sufficient for communicative success (even if it is sufficient for grasp of content). However, consider the following.

Example 2b

Stanley and Hank are at the zoo to take pictures for Hank's school project on dangerous animals. Approaching a pen, Stanley utters (2):

(2) There is a cassowary in the pen.

Stanley utters (2) with the intention that Hank believe that there is a dangerous animal in the pen that he might include in his project. Hank understands most of the utterance. He does not fully understand what 'cassowary' means; but he does know that a cassowary is a dangerous animal.

In this case, I think we should say that communication has succeeded. The illocutionary act is the same as it was in (2a); and Hank's understanding is also

²²One worry that one might have with this example pair is that the speaker's utterance plausibly has a different illocutionary force in each case – in 1b, but not 1a, the utterance is a warning. I think that a mere difference in force would not be enough to account for the difference in communicative success in these examples (e.g., the mere fact that an utterance is a warning does not obviously lessen the need to understand it). In case the reader is still concerned, the second and third example pairs do not contain this distracting feature.

²³Sociolectal theories can claim that there is more to recovering content than grasping the concepts expressed by the terms apparent from an utterance's surface form (and the manner in which they are combined). Subjects may also need to grasp the explicature of an utterance, for example.

no different – all that has changed is Stanley’s perlocutionary intention and the interlocutors’ background beliefs about the project they are engaged in.²⁴ But, intuitively, this places different demands on Hank: all Hank need understand is that there is a dangerous animal in the pen for communication to succeed relative to Stanley’s intention. In this example pair, unlike the first, neither exchange requires a greater degree of understanding from Hank than the other; rather, it is the requirements on the *dimensions* of Hank’s inferential understanding that are different. In (2a) Hank’s understanding must be such that he believes that a cassowary is a flightless bird; in (2b) his understanding must be such that he believes that a cassowary is a dangerous animal.

Example 3a

Sigfrid and Helene are planning an itinerary for their holiday in Poland. Helene has earlier expressed an interest in visiting sites relating to Polish music history. Sigfrid utters (3) to Helene with the perlocutionary intention that Helene believe that there’s a statue of a famous Polish musician in Warsaw so that she might add this to the itinerary.

(3) There’s a statue of Paderewski in Warsaw.

Helene understands Sigfrid’s utterance quite well. However, although she knows a lot about Paderewski’s political career, she is unaware that he was also a musician.

In this example, as in (2a), it is reasonable to stipulate that Helene recovers the right content. Also, once again, it is important to stipulate that this is not a case in which the hearer infers new information about Paderewski from the exchange. We can imagine that Helene is just unsure as to why Sigfrid is giving her this information. In this example, then, Helene grasps the correct content, and she even (inferentially) understands it rather well – knowing as she does that Paderewski was a Polish politician. Nonetheless, relative to Sigfrid’s intention, communication fails. Now compare this with (3b).

Example 3b

Sigfrid and Helene are planning an itinerary for their holiday in Poland. Helene has earlier expressed an interest in visiting sites relating to Polish political history. Sigfrid utters (3) to Helene with the intention that Helene form the belief that there is a statue of a famous Polish politician in Warsaw so that she might add this to the itinerary.

(3) There’s a statue of Paderewski in Warsaw.

Helene understands Sigfrid’s utterance quite well and knows that Paderewski was a famous Polish politician.

²⁴One might object that the change in background beliefs of the hearer affects her understanding of the content she recovers (one might think this if one was a holist about inferential understanding, for example). Thus, there has been a change in understanding across the example pairs. Fortunately, I think this interpretation of the examples is consistent with my account: the kind of understanding required for success is still different across the two examples (despite no change in the thought content expressed). And, moreover, the kind of understanding required appears to track the perlocutionary intentions of the speaker: it is not the case that the hearer simply needs better understanding for communication to succeed – she needs to understand the content along different dimensions depending on the speaker’s perlocutionary intentions. The speaker’s perlocutionary intentions may also determine which background beliefs the hearer must have regarding the project they are engaged in.

In this example, I think we should say that communication has succeeded. As with the previous two examples, Helene's understanding of (3) is no different than it was in (3a). However, Sigfrid's perlocutionary intention (as well as the interlocutors' background beliefs about the project they are engaged in) has changed and, intuitively, this places different demands on Helene. As with the previous example pair, neither exchange requires *more* understanding from Helene than the other; rather, the dimensions along which Helene must understand the content expressed by the speaker are simply different.

What these pairs of examples aim to show is that there is at least a kind of communication that is not captured on an invariantist approach. The examples motivate the claim that, for some cases, communicative success is context-dependent in the following way: holding fixed the thought content expressed by the speaker, the standards for successful communication can vary with features of the broad context of the speech exchange. Specifically, these features are the perlocutionary intentions of the speaker, in concert with the interlocutors' background beliefs about the purpose of the exchange. Both *how well* and *in what respects* the hearer needs to understand the speaker's utterance depends on these features of context.²⁵ Communication succeeds when (or to the degree that) the hearer's understanding is correct along the dimensions relevant to the speaker's perlocutionary intentions in relation to the project in which the interlocutors are engaged.²⁶ Having motivated this context-dependent approach, I will now go on to consider how an objector might reinterpret the examples that I have given in support of it.

5 Communicative Success and Intuitions

In each example pair, I claim, the first example is a case of communication failure and the second is a case of communicative success. In my presentation of examples, I was relying on intuitive judgments. Is there anything more that can

²⁵This approach is similar to that argued for in Bezuidenhout (1997). Bezuidenhout's primary aim in that paper is to argue that the communication of *de re* thoughts does not require shared content. She argues for a similar content account. However, in the course of arguing for the position, she suggests a context-dependent version of it. Her account is different from mine in that she thinks that preservation of reference is required for success, and she does not appeal to perlocutionary intentions.

²⁶One might object that this account is simply a different kind of invariantism: one that requires that the hearer grasp the speaker's perlocutionary intention. However, I think this claim would be an *addition* to the account I am offering here rather than an alternative. I do not claim that the hearer must represent the speaker's perlocutionary intention (nor even her illocutionary intention). My claim is that success requires only that the hearer understand the content she recovers in a particular way. If we add to the account that she must also grasp the speaker's perlocutionary intention, it will still be true that success requires that her understanding of the illocutionary content expressed by the speaker is different in different contexts. This response, I think, will also be relevant to those who would claim that what the hearer must grasp is an implicature (for those who think that implicatures can be entailments – e.g., Bach 2006).

be said to persuade someone who simply lacks these intuitions? When it comes to communicative success, it is difficult to say too much in support of an adjudication of success or failure without begging the question in favour of one's preferred theory. For example, it would not be convincing to merely insist that communication succeeds, in example 2b (but not 2a) because Hank possesses the dimensions of understanding determined to be necessary by Stanley's perlocutionary intention – that is no more than a statement of the theory that I am trying to motivate. This cuts both ways, however. It would also not be enough to claim that communication succeeds (in both 2a and 2b) because Hank grasps the correct content (or knowingly grasps it, etc.). Fortunately, I think that there is a little more that can be said in favour of the context-dependent approach. The intuitions that I am relying on will, I think, be reflected in the practices of language users in general (and this should be treated as a mark in favour of the approach).

In our day-to-day interactions, we frequently take ourselves to have uncovered misunderstandings. That is, we have an informal practice of judging whether exchanges have been successful. Pagin (2008) describes this practice as follows:

What we have is a common sense practice of judging communicative success or failure in vernacular terms such as 'He did not understand,' 'She misinterpreted him,' 'He got the message,' etc. As speakers of English we do have intuitions about the correctness of such utterances with respect to various scenarios, even if we have much poorer intuitions about the exact content of the judgments expressed, about the evidence that is appropriate for making them, or about the extent to which the judgments are supported by the evidence used. (Pagin 2008: 86)

Furthermore, when we do believe that a misunderstanding has occurred, there is characteristic behaviour that we engage in to correct this perceived misunderstanding. Rapaport (2003) calls this behaviour 'negotiation': it is the process by which understandings are minimised to the point at which they are judged small enough to be ignored.²⁷ Hearers, if invested in the exchange, will say things like 'I don't understand you' or ask 'What did you mean by that?' And speakers, if cooperative, will offer more information. Pagin suggests that an approach to communicative success should offer diagnoses of success and failure that more or less coincide with the judgments of this common-sense practice. He takes this to be an adequacy criterion on a theory of communicative success; and I think, at the very least, it should be treated as a strong mark in favour of an approach. It would be strange – and in urgent need of explanation – if one's theory predicts that most human beings are usually very wrong in their judgments of communicative success and failure.²⁸ So, how does the context-dependent account fare in this regard? Consider example (2a) again. It is likely that someone who found themselves in Hank's position

²⁷Negotiating behaviours are also discussed in Clark 1996.

²⁸Humans can make mistakes and lack information, of course. The idea is that, if provided with relevant information, a typical speaker's judgments as to success and failure of an exchange are likely to be roughly accurate.

would engage in the typical negotiations of a hearer who takes themselves to lack understanding. While Hank might come to form the belief that there's a cassowary in the pen, he will be confused as to why Stanley made this assertion. He might seek further clarification, asking, for example, 'Why are you telling me this?' Even if he does not voice his confusion in the moment, he may continue to wonder until he discovers the reason that Stanley said what he did. Similarly, speakers also engage in negotiations: if Stanley finds out about Hank's partial understanding in (2a) – as he might if Hank voices his confusion – he will likely explain that a cassowary is a kind of flightless bird and thus relevant to Hank's interests.

But what of Hannah in example (1b)? An objector might point out that Hannah will attempt negotiations in this case too, despite my claim that communication has succeeded here. For example, it wouldn't be surprising if Hannah were to utter 'Who are you pointing at?' in response to Sally's utterance, even when she does understand well enough to take steps to protect her health. Isn't that – by my lights – a reason to think that communication has failed after all? I think that there is something to this complaint, but it is something that my account can accommodate. What is important, for my purposes, is that there is *some* sense in which communication succeeds in (1b). My claim is that communication succeeds *relative to Sally's perlocutionary intention that Hannah avoid getting sick*. However, because we are considering success relative to intentions in this way, I think I can simultaneously accommodate the thought that (depending on how we fill out the details of the example) there is some further sense in which communication fails. The reason for this is that speakers can have more than one perlocutionary intention when communicating, and this can place multiple demands on the hearer. Imagine that, in (1b), Sally has two perlocutionary intentions: that Hannah believe that Bill is eating shellfish *and* that Hannah avoid getting sick. If this is the case, then, according to my account, communication will have *succeeded* relative to one of Sally's intentions and *failed* relative to another. Which intentions Sally has (and thus the standards for success) will, I think, often be reflected in the interlocutors' negotiations. For example, suppose Hannah does ask 'Who are you pointing at?' There are a number of ways in which Sally could reply that would indicate one or more different perlocutionary intentions on her part:

- (4) It doesn't matter! Just get out of here.
- (5) I was pointing at Bill. Shouldn't you get out of here?
- (6) Bill – I thought he was a vegetarian!

What negotiations take place (and when they terminate) will depend on what the speaker was trying to do, and this is better accommodated on my context-dependent approach. My point here is not that interlocutors' negotiations will always be a perfect guide to a speaker's original intentions; as Clark (1996: 195) notes, speakers will sometimes revise their intentions as a conversation progresses. Moreover, sometimes communication failures will never be discovered. Rather, I think it is a mark in favour of my view that it can explain why a speaker might react in these different ways to confusion on the part of the hearer. An invariantist does not offer an explanation of this – or, rather, she must claim that these different reactions are not

relevant to the question of whether communication has succeeded. One advantage of my proposal, then, is that it can accommodate the fact that negotiations may cease even when the content grasped by the hearer is not very well understood and, similarly, why negotiations may sometimes continue even beyond a very high degree of understanding.

In the preceding, I don't claim to have shown that the sorts of example under consideration will be widespread. How often is communicative success context-dependent in the way I am proposing? I think there is reason to think that examples like the ones above will be quite common. Returning to the earlier thought from Clark, communication (of the sort under consideration here) is a goal-directed practice that we engage in to coordinate action in cooperative projects beyond the communicative exchange itself (for example, completing a school project, or planning a holiday). Utterances, on this approach, are coordination devices, instrumental in securing further goals. I think we should expect to find examples like the ones presented above whenever the same utterance (expressing the same thought content) can be used as a coordination device in different cooperative projects. Because, as we have seen, a hearer's understanding of an utterance or thought content can vary along different dimensions, utterances will very often be multi-purpose devices. That is, for each dimension along which an utterance content can be understood, there is (at least) one perlocutionary effect which, under the right circumstances, that utterance can be instrumental in securing. Successful communication is a matter of understanding the utterance content along these contextually-determined dimensions.

6 Conclusion

In this paper, my aim has been to motivate a context-dependent approach to communicative success. This approach treats the standards for success as context-dependent in the following sense. Holding the thought content expressed by the speaker fixed, features of context can place different demands on a hearer with respect to the manner in which she must understand the content that she recovers. I identified two features of the context that can play this role: the perlocutionary intentions of the speaker and the interlocutors' background beliefs. I don't claim that this approach captures the *only* kind of communication that interlocutors can engage in. Rather, my claim is that there is an important notion of communicative success that is relevant to the coordination of action (but not defined in terms of it), and this notion of communicative success cannot be accommodated on a purely invariantist theory.

Acknowledgements I would like to thank Lewis Pollock, Bryan Pickel, and Andrew Peet for very helpful comments on previous drafts of this paper. I also thank audiences at conferences and workshops at the University of Warsaw, University of Tartu, University of Edinburgh, and Ruhr University, Bochum.

References

- Austin, J. L. (1962). In M. Sbisà & J. O. Urmson (Eds.), *How to do things with words* (2nd ed.). Oxford: Oxford University Press, 1975.
- Bach, K. (1987). On communicative intentions: A reply to Recanati. *Mind & Language*, 2, 141–154.
- Bach, K. (2006). The top 10 misconceptions about implicature. In B. Birner & G. Ward (Eds.), *Drawing the boundaries of meaning. Neo-Gricean studies in pragmatics and semantics in honor of Laurence R. Horn*. Amsterdam: John Benjamins.
- Bezuidenhout, A. (1997). The communication of *de re* thoughts. *Noûs*, 31, 197–225.
- Bezuidenhout, A. (1998). Is verbal communication a purely preservative process. *The Philosophical Review*, 107, 261–288.
- Bird, A. (2002). Illocutionary silencing. *Pacific Philosophical Quarterly*, 83, 1–15.
- Buchanan, R. (2014). Reference, understanding, and communication. *Australasian Journal of Philosophy*, 92, 55–70.
- Burge, T. (1979). Individualism and the mental. *Midwest Studies in Philosophy*, 4, 73–121.
- Burge, T. (1993). Content preservation. *The Philosophical Review*, 102, 457–488.
- Cappelen, H., & Lepore, E. (2005). *Insensitive semantics*. Oxford: Blackwell.
- Cappelen, H., & Lepore, E. (2007). Relevance theory and shared content. In N. Burton-Roberts (Ed.), *Advances in pragmatics*. Basingstoke: Palgrave Macmillan.
- Carston, R. (2001). Explicature and semantics. In S. Davis & B. Gillon (Eds.), *Semantics: A reader*. Oxford: Oxford University Press.
- Clark, H. (1996). *Using language*. Cambridge: Cambridge University Press.
- Davidson, D. (1986). A nice derangement of epitaphs. In E. Lepore (Ed.), *Truth and interpretation: Perspectives on the philosophy of Donald Davidson*. Oxford: Blackwell.
- Dummett, M. (1978). Frege's distinction between sense and reference. In *Truth and other enigmas*. London: Duckworth.
- Evans, G. (1982). In J. McDowell (Ed.), *The varieties of reference*. Oxford: Clarendon Press.
- Fodor, J. (1994). *The elm and the expert*. Cambridge, MA: The MIT Press.
- Fricker, E. (2006). Testimony and epistemic autonomy. In J. Lackey & E. Sosa (Eds.), *The epistemology of testimony*. Oxford: Oxford University Press.
- Goldberg, S. (2007). *Anti-individualism: Mind and language, knowledge and justification*. Cambridge: Cambridge University Press.
- Grice, H. P. (1957). Meaning. *The Philosophical Review*, 66, 377–388.
- Heck, R. (1995). The sense of communication. *Mind*, 104, 79–106.
- Jorgensen, A. (2009). Holism, communication, and the emergence of public meaning: Lessons from an economic analogy. *Philosophia*, 37, 133–147.
- Langton, R. (1993). Speech acts and unspeakable acts. *Philosophy and Public Affairs*, 22, 293–330.
- Longworth, G. (2010). A plea for understanding. In S. Sawyer (Ed.), *New waves in philosophy of language*. Basingstoke: Palgrave Macmillan.
- Longworth, G. (2018). Understanding what was said. *Synthese*, 195, 815–834.
- Millikan, R. (1984). *Language, thought, and other biological categories*. Cambridge, MA: MIT Press.
- Newman, A. (2005). Two grades of internalism (pass and fail). *Philosophical Studies*, 122, 153–169.
- Pagin, P. (2008). What is communicative success? *Canadian Journal of Philosophy*, 38, 85–116.
- Pollock, J. (2015). Social externalism and the problem of communication. *Philosophical Studies*, 172, 3229–3251.
- Putnam, H. (1975). The meaning of “meaning”. In K. Gunderson (Ed.), *Language, mind, and knowledge*. Minneapolis: University of Minnesota Press.
- Rapaport, W. (2003). What did you mean by that? Misunderstanding, negotiation, and syntactic semantics. *Minds and Machines*, 13, 397–427.

- Recanati, F. (1986). On defining communicative intentions. *Mind & Language*, 1, 213–242.
- Recanati, F. (1993). *Direct reference*. Oxford: Blackwell.
- Sperber, D., & Wilson, D. (1986). *Relevance communication and cognition*. Oxford: Blackwell.
- Stalnaker, R. (2002). Common ground. *Linguistics and Philosophy*, 25, 701–721.
- Tomaseello, M. (2008). *Origins of human communication*. Cambridge, MA: MIT Press.

Truth and Context



Gerald Vision

Abstract It is widely, though not universally, acknowledged that the truth of fully determinate propositions often depend on the contexts in which they occur. In addition to that fairly innocuous claim I venture the much bolder claim that context can aid in demonstrating the truth of a number of propositions that ascribe homely properties to fictional (and, I assume, non-existent) subjects. After summarizing salient and what I deem to be worthwhile objections to my view, I set out to expose their flaws. With that in the archives I build the case for my conclusion. It consists in showing how accounting for differences among cases make judgments of truth and falsity unavoidable, and how to circumvent some familiar pitfalls.

Keywords Truth · Context · Semantics · Existence · Reference

The importance of context as one determinant of truth is generally acknowledged. Thus, various sets of philosophers have emphasized not only the role of context in specifying indexicals, demonstratives, and other deictic elements, but also its importance in determining the truth of conditionals, modals, knowledge ascriptions, and for truthbearers whose predicates are gradable, vague, or normative. However, I maintain that we should boldly take this a step further, to statements made with sentences whose subjects are the names of fictional entities, and that this can be accomplished without the usual ‘in fiction’ qualifiers or other substantial reconfigurations of the propositions expressed. On the view I support, context is enough to bring off their truth candidacy. To make that case we should include predicates that do not fall in the classes to which I have just alluded, predicates such as ‘danced’, ‘hunted’, or ‘wears a fedora’. However, for simplicity I shall focus on properties picked out by commonplace sortals such as ‘is a detective’, ‘is a horse’, ‘is a widow’, and ‘is a soldier’. To be specific my contention is that in the right

G. Vision (✉)
Temple University, Philadelphia, PA, USA
e-mail: gvision@temple.edu

context stated sentences such as ‘Indiana Jones is an archeologist’ may be true without further (implicit) adornment. By the end of the present discussion I hope that this very audacious proposal will look more like the merest commonsense.

1

First some brief comments on the title’s first conjunct, truth. I assume a loose-limbed Correspondence Theory. It is the default position for truth theory. *Qua* default, past clarifications the case for it rests largely on its defense against attacks, including those that question its capacity for supplying details. As such it is worthwhile to pursue further reflections on its assumption.

In calling this view “loose-limbed” I have in mind correspondence with states of affairs (or facts, situations, events, thick individuals, or whatnot) forming the core of our understanding of truth’s conditions of satisfaction. We can remain neutral on the admissibility of controversial cases that may bear various sorts of semantic or ontological dependence on that core (so-called secondary cases). For starters, we can work from the following schema

(C) A proposition p being true is constituted by its correspondence to a state of affairs Σ , *ceteris paribus*.

By ‘proposition’ I mean only the content of an assertion or a propositional attitude. In particular, I do not identify, individuate, or define propositions as sets of possible worlds, although propositions can be so represented once they are at hand. Propositions of interest are the structured upshots of speech or thought episodes, and are analyzable into their elements, whether those are concepts, referents, or their syncategorematic accompaniments. They are generated by concrete episodes, and may be extended to a proper subset of realistically possible concrete scenarios, although still as structured combinations. On this view there were no propositions—say, in the Precambrian era—before there were cognizing agents. On the other hand, non-linguistic *facts* were present from the beginning of time. But propositions came into existence no earlier than individuals who could think or state them.

The notion of a proposition has been valued in philosophy as a way of introducing an exclusive, primary, or reliable bearer of truth and falsity.¹ This is due in large part to a proposition’s ability to account for the fact that different sentences may be used to make the same assertions and by a capacity to appear in larger compounds. Because propositions are upshots of actual occasions, for ease of exposition when considering uncompounded sentences I shall also use *statements* interchangeably.

These are points to store in the background. The rest of the discussion proceeds as follows. §2 is devoted to explaining the use I plan to make of context. §3 explores what I take to be an indisputable insight from Tim Crane’s book *The Objects of*

¹See Richard Cartwright (1987).

Thought and his follow-through, at which point I dissent from his view that fictional utterances such as

(S) Sherlock Holmes is a detective

are always false. Tenses aside, I shall take this instance as a generic stand-in for fictional sentences of subject-predicate form. §4 lists the subsequent reasons Crane gives for that claim, and in §5 I present my grounds for rejecting his conclusion. §6 is devoted to arguments for extending the power of context to support truth even for assertions whose topics are, as in (S), fictional objects and happenings.

2

There is an abundance of commonplace evidence suggesting that the truth-value of a significant number of *fully-determinate propositions* rests on the *contexts* in which they are uttered. Let us call that view, which I shall defend below, **T-Contextualism**. The view needn't hold that context is a factor in every truth. Perhaps, arithmetical truths, that the Sun is larger than the Earth, and many others are true in context zero. But as J. L. Austin writes (taking *stating truly* for *stating what is true*)

[I]n the case of stating truly or falsely, just as in the case of advising well or badly, the intentions and purposes of the utterance and its context are important; what is judged true in a school book may not be so judged in a work of historical research (1962: 142).²

The title 'Contextualism' has been used for a number of different theories in roughly the same area.³ However, the contexts I have in mind are only those for the truth (or falsity) of determinately fixed propositions. It is the world, not the truthbearer that is contextually relevant. In particular it is important to distinguish the present thesis from two methods for completing the bearers of truth with which it may be confused.

To begin, truth-conditions have been cited as playing a semantic role, determining the *contents* of the sentences for which they are conditions—if not strictly determining meaning, at least determining whatever contributions pragmatics may make to that content. On that view truthbearers have not been finalized independently of their truth-conditions. But T-Contextualism concerns only *truthmakers*, the worldly correlates or possibilities relevant to evaluating truth, not contributions to the propositions or sentences under evaluation. In fact, truth-conditions have not always been viewed as involving nothing beyond semantics. Thus, if one holds that *truth-conditions* can change with context of utterance, but without altering the truthbearer,

²J.L.A. is careful not to say that these *are* true, but the suggestion is unmistakable. Here I am defending the truth of the statements, not merely our entitlement to judge so.

³For example, François Recanati writes "'Contextualism' is my name for the doctrine that *no* proposition could be expressed independent of context" (1993: 267n5, my emphasis). I have no objection to the view or its use of the label, but it is not the view now being considered.

they are eligible to figure in T-Contextualism. John Searle's view discussed at the outset of §6 appears to be a thesis of this sort. But care is needed to avoid the re-interpretations of those who would absorb any contextual features into an enlarged semantics in defense of their 'nothing beyond semantics' view of truth-conditions.

As construed here T-Contextualism has no implications for ways to complete a proposition by bringing out unexpressed constituents. Two sorts of supplements are especially prominent. First, *saturation* specifies the referents of indexicals, demonstratives, quantifiers, and tenses to complete the statement made. Next, what Recanati (1993, 2005) calls "[f]ree enrichment—the process responsible for making the interpretation of an utterance more specific than its literal interpretation" (2005: 177) plays a role. In Geoffrey Nunberg's (1979) frequently cited example, one waiter says to another "the ham sandwich didn't leave a tip".⁴ These are not cases of T-Contextualism.

T-Contextualism concerns only the states of affairs relevant to, and thus eligible for, the correspondent side of the equation. Put otherwise, it is a species of a *relevant alternative* view (RA for short), invoking a need for a certain dimension of evaluation's circumstances.⁵ Austin has pointed us toward a host of potential examples. Another might go as follows. Maud is seated on the front passenger seat of a parked car whose door is open while one of her feet is planted on the pavement. Someone in the general vicinity looking for Maud might be told that she is in the car. On the other hand, a passenger in the back seat might in truth warn the driver not to start the car because Maud isn't in it yet. My contention is that both remarks can be simultaneously true in the cited circumstances.

For one more example to stir the pot, and taking inspiration from a case devised by Dan Sperber and Deidre Wilson (1985/1986), if talking to a friend in foreign parts about Al's residence it may be true to say that he lives in Chicago, although he recently moved to the suburban side of the street at the city limits, whereas it would be false to utter those words when vouching for Al's residence to an official at a local election.

To the best of my knowledge there is no decisive refutation of the counterclaim that in those cases there is always a further adjustment to the truthbearer. For example, someone might venture that "[t]he whole utterance—the context and words uttered—is relevant to identifying the unarticulated constituent",⁶ even if further details are not forthcoming. That view would rule out, a priori, any T-Contextualist truthbearer. In spite of its coherence, appearances run strongly against such a semantics, seemingly broadened only to place every case under the umbrella

⁴For a number of additional examples see Robin Carston (1988).

⁵It is loosely fashioned after, and analogous to, the RA view for the conditions of knowledge pioneered by Fred Dretske (1970, 1972) and those for truth sketched by David Lewis (1979). The central objection to Dretske's view regarding knowledge closure raises no analogous threat to the current view.

⁶Crimmins and Perry (1989: 700).

of semantics. Austin's description of those sorts of occurrences is much less contrived.

The following two provisos seem to be central to T-Contextualism:

- (P1) context is a determinant of the truth-value of at least some statements,⁷
 (P2) a difference of context can make a difference to a statement's truth-value.

We may illustrate T-Contextualism with statements made with simple sentences of the form ϕa . But because my aim is eventually to discuss sentences from fiction, let us first consider their ostensible referential components that, on just about anyone's account, fail.

3

Following Tim Crane's lead I begin from instances of 'thinking about'. This gives primacy to *persons* as subjects of 'proto-referring' rather than to that of *terms*. Crane is forthright about the fact that our thoughts can be about nonexistent objects.⁸ He resists framing his *talking about* in vocabulary from the 'reference' family,⁹ but it is a clear and natural way to put things. The central point is that under either heading the notion looks to be beyond dispute. Although we are clearly aware of Sherlock Holmes' fictionality, how can it be plausibly denied that *we* are thinking about Holmes? Similarly for language use as well as thought, *we* can refer to Holmes. Emphatically, this is not because non-existent, fictional characters have a more exotic kind of being: no spurious ontic status in addition to plain existence is being invoked.

Certain critics have considered reference based on notions such as talking or thinking about to be too feral for systematic theorizing. Nevertheless, if our interest does not spring from just such notions, however it is subsequently refined, the central concern of philosophers with referential practice (that is, with the relation of language/thought to the non-linguistic world) is a mystery to me, and it is simply brutally obvious that we engage in thought and language in this way quite apart from any further finessing of ontology.

⁷I say 'a' rather than 'the' determinant because fit with the world, the rough correspondence component, is always a factor in the relevant cases.

⁸Future page references to Crane are to his 2013. He rejects talk of non-existent *entities* (5), but is content to write about non-existent *objects* (of thought and talk) (16), *things* (12), or even *items* (27).

⁹The one reason he gives is that it conforms to "the usual philosophers' practice of calling names like 'Pegasus' *non-referring*" (46). However, that practice is far from universal, and concerns the adoption of what is to be called 'term reference' as the only legitimate such notion. Moreover, it is at odds with Crane's own insight. I shall continue to talk here about personal and term reference. (A brief excuse for this appears in §6.) My choice of terms will make no difference to the substance of Crane's views.

On occasion attempts have been made to analyze away that sort of engagement. I cannot go into detail here, so I'll satisfy myself with the (dogmatic) claim that in the past those who have wished to dispatch the data by various methods have taken on more implausible commitments. Indeed, it is difficult to take seriously any view that denies that *we* think and talk about, thereby refer to, what doesn't exist. How could we investigate a hypothesis about the existence of a shadowy something such as Yeti, Agamemnon, The Templars' hidden treasure, or a minor character in the Bible if it turns out that the very form and proper understanding of the question being raised depended on an analysis that in turn depended on the correct answer to its inquiry? Doesn't this get things backwards?

A number of further issues raised by Crane's account are bypassed. Here I confine myself to the interplay between what seems to me his incontestable starting point and statements he then makes about certain truth-evaluations involving fictional objects.

If *we* can refer to Sherlock Holmes, a natural next step would seem to be that we can frame truths about him. However, Crane holds that

(S) Sherlock Holmes is a detective

is false (23).¹⁰ One might allege various other truths about Holmes such as that he lives in Baker Street, plays the violin, uses cocaine, is brilliant, and is an expert on tobacco stains, each of which fail for one reason or another. However, for Crane the problem is with the properties predicated. He calls such workaday properties 'existence-entailing' (2013: 59); they cannot be truly predicated of a non-existent object.¹¹ "Objects have natures. What their natures are is a matter of empirical or metaphysical study" (63). Fictional objects have no natures to study further in either way. (This might be strictly consistent with T-Contextualism if we regard the one and only context for those sorts of predications as actuality as opposed to nonexistence. Then condition (2) of T-Contextualism simply need not come into play.) Crane does acknowledge other truths about Holmes that are not existence-entailing, such as that he is a fictional character, is widely admired, or is thought about.

As advertised earlier, I later defend the opposing view that examples essentially like (S) can be true when the context provides the appropriate background against which to assess it. But before offering my counterproposal much preparatory spade work awaits. First, three important preliminaries must be mentioned. Next I believe

¹⁰I distinguish these cases from those of false scientific, theological, etc. hypotheses, which in many instances warrant a different sort of treatment. My remarks are not intended to form a perfectly general account about reference to whatever non-existent one can dream up.

¹¹Two points. First, for simplicity I assume that all predications express properties, even 'exists' and 'does not exist'. This does not affect the points under discussion. Next, Zalta (2000) claims that existing subjects *exemplify* their properties, fictional ones *encode* them; for Van Inwagen (2000) existing subjects *have* their properties and fictional ones (merely) *hold* them. So far as I can tell these are distinctions in search of a difference. Also, it is unclear how those property distinctions change the sentences in which they appear. Like Crane, I ignore those distinctions here.

it will be a useful object lesson to examine Crane's arguments, in §§4-5, for declaring (S) to be false.

First preliminary, the distinction Crane draws, if not in the terms he prefers, boils down to a familiar dichotomy between *personal reference*, references as made by cognitive subjects, and *term reference*, the reference of the expressions themselves.¹² Personal reference licenses one to think and speak about both existents and non-existents. On the other side, Crane calls term reference a technical notion (9), governed in part by the requirement that referents must exist.¹³ Ordinarily we needn't distinguish the two because in a typical run of cases personal references are also to existents. But when the two notions come apart in a particular case it produces controversy and sets in motion theory-driven crisis management. "If term rather than personal reference is what matters for truth, what are we to say about the lack of an existing referent?" The solution frequently adopted has been to deny that a proposition has been expressed. Those who allow that a proposition has been expressed may hold that it lacks truth-value or is simply untrue (but not false). Yet others claim that the proposition is false. Crane chooses falsehood, perhaps because he locates (S)'s failure in the existence-entailing nature of the predicate rather than in the proposition's subject.

Next, negative existentials that appear to the philosophically uncorrupted to be obvious truths, such as 'Vulcan does not exist' and 'The round square does not exist', present a problem for many philosophers because a condition for their truth is the existence of their respective referents. Crane could have a simple solution to those worries—although not the one he chooses—because personal reference has no such condition and for him ordinary quantification lacks existential import. (See his Chapter 2, especially 40.) Those two theses should also take off the table one argument, albeit not Crane's, for preventing (S)'s truth. Thus, when talking specifically about those parts of Crane's view that I have claimed to be beyond dispute, I omit an oftentimes expected discussion of this familiar conundrum.

Finally, I set aside the following three leading strategies for defending (S)'s acceptability. First, some regard propositions about goings-on in fiction to be understood as governed by a prefix along the lines of 'in such-and-such a fiction'. Ordinarily that prefix will be unexpressed, but its insertion would be a way to turn certain failed propositions into truths. Second, one may take fictional figures to be abstract entities and thereby suitable existents for reference and ordinary predications. Finally, one may regard all fictional utterances as occurring inside a setting of make-believe, in which some pretended utterances are preferred. It would

¹²Cf. Donnellan (1966), Kripke (1977). These types have also gone under the labels 'speaker's reference' and 'semantic reference', respectively. Crane's misgiving about this terminology is discussed in note 9.

¹³Of course, 'technical' needn't mean invented whole cloth; it can have a firm basis in ordinary discourse, regimented only to keep that discourse from going off the rails. Nevertheless, the sense in which a condition promoted by Plato in the *Sophist*, and by Parmenides before him, is technical is not as clear as one would hope. Did they suppose they were imposing a specialized vocabulary?

take me too far afield to explain satisfactorily why I reject each strategy.¹⁴ Crane also explicitly rejects the first two and his declaration that (S) is false is incompatible with accepting that it is only a pretended proposition. But more importantly, my later defense of (S)'s truth will not appeal to any such recommendations for reconfiguring its utterance.

4

Let's look at the reasons behind Crane's claim that (S) is false. Taken individually they are admirably clear, but what counts as a distinct argument in his discursive exposition is less transparent. So I merely list those explicit and implicit points of his critique that strike me as most central to his case.

Of course the existence-entailment of ordinary properties stands out. Aside from metaphysics and other a priori subjects, having a nature implies having a number of further spatiotemporal and natural properties open to empirical inquiry. For example, if Holmes is a detective he would have parents and would occupy a specific location at times other than those implied in an Arthur Conan Doyle opus. But because fictional characters lack natures they are not proper subjects for an investigation into those particulars. Author's sayso ends the matter.

Next, aside from the specifics of the existence-entailment thesis, these predicates can be the source of unwelcome implications. Compare (S) with a statement about an actual Scotland Yard agent,

(I) Sir Ian Blair is a detective.

Is it not something on the order of a category mistake to say in a single breath that Sir Ian Blair and Sherlock Holmes are detectives?

In addition, when inquiring into Holmes' properties we encounter a number of indeterminacies: once we have concluded that a multitude of particular incidents of Holmes' life are not implied in Doyle's tales, it makes no sense to inquire about them. For example, although, we may infer that Holmes' Baker Street residence was just down the street from Regents Park, does it make sense to ask whether he ever strolled in the park? If he were an actual person, that question would make sense even if we were unable to answer it.

¹⁴I offer reasons for rejecting the preface strategy in Vision (1993). The gravamen of my charge there is that the intensional contexts created by fictive prefaces prevents us from using them to draw inferences in conjunction with even the most innocuous commonplaces. Treating fictive sentences as detached bulletins does scant justice to anyone's reasons for wanting to preserve them. As for abstractions, they could not possess most of the properties of fictional characters we want to account for. And whereas the last strategy, inspired by Kendall Walton's (1990) make-believe account, may withstand scrutiny in the simple cases with which Walton illustrates it, attempts to extend it from authorial to non-authorial utterances stretches the notion of pretense to extremes that I believe only those with prior leanings in that direction would be inclined to countenance.

Beyond that, in elaborating the distinction between properties that are and are not existence-entailing Crane makes intermediate use of a distinction, taken from Stephen Schiffer (2003), between *substantial* and *pleonastic* properties. Pleonasms lead to transformations that get what Schiffer (2003: 51) describes as “something from nothing”. Crane disparages ordinary (roughly, non-intentional) properties predicated of fictional figures as just such pleonasms.

Pleonastic properties are further elaborated as “representation-dependent” (68), a term of art borrowed from Colin McGinn (2000: 37). This is the case for the property of being a *fictional* anything and also for stating, say, that Holmes is more famous than Sir Ian Blair because it is “dependent on the fact that more people are familiar with Holmes than they are with Sir Ian” (68–69). If the subject should even happen to exist its representation-dependent properties never imply that it exists in the way that ‘is a detective’ is presumed to do.

I call the pleonastic-substantial distinction ‘intermediate’ because it is introduced mainly to flesh out Crane’s predilection for a broader minimalism in which a minimalist theory of *truth* is central. On that view truth can be summarized (at least for uncontroversial cases) by the schema

(D) The proposition that *p* is true if and only if *p*.

No doubt, truth theorists of every stripe can accept (D), but for minimalists (D) itself is a complete theory for uncontroversial cases. Truth “does not have a hidden nature, . . . nothing for metaphysics or science to investigate” (65). To contrast with my opening salvo on behalf of (C), this surely implicates the rejection of *fact* and a *correspondence* relation for a theory of truth. An apparent advantage of this minimalism is avoiding getting bogged down in the messy business of truthmaker theory. Crane clearly implies that much.

There is a sense in which these objections are just right for the occasion. We could have started from difficulties arising from a misfit with formal constructions. However, it seems best to first find problems inherent our actual practices to bolster the *bona fides* of reforming common discourse or for discovering its hitherto undetected underlying forms. Moreover, without the failed analyses mentioned in note 14 coming to (S)’s rescue, the case against (S)’s truth is impressive and is likely to be found among considerations of this order. If the bald truth of (S) is to be defended, it must address arguments like the foregoing. I now turn to that task.

5

The first objection denies that Sherlock Holmes has existence-entailing properties. If he is a detective he would have many other properties of existing detectives. He must have a heart, a liver, a spine, and be located somewhere; but, Crane responds, how can a non-spatiotemporal thing such as Holmes have those features? In reply I maintain that Holmes has these properties in just the same way that he is a detective. His regular residence is a flat on Baker Street, he has arms, legs, and lungs, sleeps,

eats, time passes for him, and so on. *Perhaps* these properties are also pleonastic, but that is to be expected if Holmes' status as a detective is judged so as well. (More on that presently.) Much of fiction is founded on what readers can expect without having their sensibilities jarred. Although we may never have a reason to state that Holmes has a heart, given that he is presented as a specimen of homo sapiens that is surely implied in the author's writing.¹⁵ Taking into account the license of authors for extreme science fiction and cartooning, any inference we make on this basis is, in Grice's terms (1989: 44–46), 'cancelable'.¹⁶ But until actually canceled these are legitimate and reasonable implications. A merely cancelable implication is an implication unless canceled. There may be further reasons for rejecting these ascriptions, but not due to a lack of those particular properties or a failure to conform to normal predicational practice. If an author credits Holmes with being a detective, the author and us are not entitled to withhold the usual accompanying properties without giving readers explicit directions to the contrary. Ignoring that clear directive is the only argument I can discover in Crane for such predicates being existence-entailing. In lieu of further defense, the existence-entailment requirement fails.

Certainly Holmes is also indeterminate in many respects. It is open to controversy whether existing objects are fully determinate, but without broaching that issue we should certainly acknowledge that the indeterminacy of fictional characters is much more extensive and diverse than any that might attach to existents. There is no answer to whether the number of hairs on Holmes's head is odd or even, or if he has a mole on his back. Borderline cases aside, fictional characters are indeterminate in dimensions inapplicable to their actual counterparts. Indeterminacy or its expansiveness might yield reason enough to distinguish the actual from the fictional, but that is not the same as a distinction between having and lacking all the properties that there is evidence for ascribing. It doesn't follow from indeterminacy in other respects that we have a reason to curtail those particular properties for which that larger issue has not been specifically raised. Even without Crane's envisioned empirical inquiry our judgments about ascribing or withholding a particular property of a fictional object is not normally arbitrary. Possession is often dictated by solid textual evidence and its universally presumed accompaniments.

Next, Crane takes the *representation-dependence* of a property as divorced from the existence of its possessor. However, artifacts are inherently representation-dependent, and all of their artifactual properties—including their descriptions as artifacts—depend on their being represented. Nevertheless, discounting extreme

¹⁵Cf. Urmson: "if the story begins with the words 'Tom was a middle-aged man from Columbus, Ohio,' we may assume that he is visible to other men, needs food and drink, speaks English, and so on, unless we are explicitly warned to the contrary. Without some such presuppositions as these the story will be unintelligible" (1976: 153–54).

¹⁶Crane remarks that unlike horses as members of the animal kingdom, Pegasus has no parents (62). ('He' sprang from the blood of Medusa.) That may be a good cladistic reason for Pegasus not being a horse, but as far as intrinsic features go, and whatever is implied by the addition of wings, is there any reason to suppose that Pegasus is not otherwise equine-like?

metaphysical eliminativism, they exist.¹⁷ It is patently implausible to deny that computers, radios, statues, and saucepans exist. Something wouldn't be a Rolls-Royce unless its creators represented it as such. If through a cosmically odd coincidence of bits of matter converging via a hurricane, birds dropping more parts into it from their bills, and the like, these all came together to look and work just like a Rolls-Royce, it wouldn't be a Rolls-Royce, or even an automobile. It might become an automobile if someone began driving it, and representing it as such. But *qua* original bit of matter it is not an automobile, and it is *never* a Rolls. An automobile is an artifact, and an artifact is defined in terms of its function for which a creature designs or adopts it—that is, represents it. This Rolls simulacrum was not created in that way. But to say of the Queen's auto that it is a Rolls-Royce is representation-dependent, and if any property is existence-entailing (a moot point) that one is.

In another objection Crane asserts that saying in a single breath that Sherlock Holmes and Sir Ian Blair are detectives would be wrong. However that is only because it would conflate incompatible contexts. If someone is acquainted with the name 'Sherlock Holmes' only by having overheard a conversation about a fictional character, in which Holmes' fictiivty wasn't raised, and I proclaim in a later conversation that Holmes was a detective, I may be informing my interlocutor, *inter alia*, that he was known for that rather than for being a barber, a barrister, a pirate, or anything else in that range. The only relevant context is what our knowledge of him depends on, not his existence, which was never in question. In sum, Crane's argument works only if the exclusive context for judging truth is that in which the fictional sentence is compared with actuality. Crane is not alone in resisting the intuitive pull of statements such as (S). But then it is relatively rare to discover a theorist who even addresses T-Conventionalism's challenge to this sort of objection.

What should we say about Schiffer's notion of pleonasm to deflate the use of 'is a detective' in (S)? To begin with, its invocation by Crane is problematic. He must decide whether pleonastic properties are or are not really properties. Either choice subverts his ends. First suppose that they are degraded properties (as Schiffer claims). Their being pleonastic is due to the bases on which they are ascribed, not on having different *senses* from their substantial counterparts. It would be remarkable if there were a difference in sense between 'red' predicated of a real ruby and of a fictional ruby. How could an author expect us to understand her words? Applying this to 'is a detective', if Holmes has that property pleonastically, he is a detective and (S) is true.

Now examine the alternative: pleonasms are *ersatz* properties. They are no more properties than model railroads are railroads. A difficulty with that answer is a

¹⁷*Pace* the historian Yuval Harari (2014), who writes, working on a highly refined notion of existence, "if you examine any large-scale human cooperation, you will always find that it is based on some fiction like the nation, like money, like human rights. These are all things that do not exist objectively, but they exist only in the stories that we tell and that we spread around." Perhaps human rights are mere ideals, but nations and money?

further distinction between pleonastic properties and something falling short even of being pleonastic. Holmes may be a detective only pleonastically, but he is not even a pleonastic pirate. Pleonasms allow us to make that distinction. Indeed, for Schiffer they result via inferences from truths. Rather than simply dismissing abruptly all such cases, aren't pleonastic properties introduced in the context because, say, (S) has something to it that 'Holmes is a pirate' lacks? So exactly how does Crane's notion fit?

Finally, Crane uses the pleonastic-substantial distinction to underwrite his minimalism about truth. An instance of the minimalist's (D) relevant for our purposes is

The statement (/proposition) that Sherlock Holmes is a detective is true if and only if Sherlock Holmes is a detective.

This is all the minimalist allows to truth's ordinary predication. Given that Crane grants (S) propositional status, on his own account the truth of (S) should follow from that biconditional; its left-hand side adds nothing of substance to the right side. If the predication of truth is that insubstantial it is unclear Crane can square it with his contention that (S) is false. Moreover, if he claims it is false because Holmes doesn't exist or because being a detective is existence-entailing he is offering a metaphysical grounds for its evaluation, thereby violating minimalist restrictions.¹⁸ Minimalism requires that the conditions for applying a truth predicate are little more than syntactical, the only positive conditions for a proposition to take a truth predicate being a capacity to occur in compounds and to take negation. (Restrictions may also be in place to avoid paradox.) Falsifying on grounds of lacking an existence-entailing property are precisely the explanations loved by truth realists and disavowed by minimalists.

So much for *ad hominem*. On the other side my tolerant correspondence theory requires truthmakers in addition to what is affirmed in (D). I cannot mount an adequate case against deflationism in this place,¹⁹ but it is not needed. The pertinent question is only why one should take (S) to be true on anyone's notion of truth. My simple answer appeals to T-Contextualism. Holmes being a detective accurately distinguishes this fictional character's details from countless other potential ascriptions that are off-target.

Upon review Crane's best case for rejecting (S)'s truth fails. I use him as target only because I deem his beginning with what I am calling 'personal reference' as invaluable—indeed, mandatory—and its not being carried through for truth leads us into a quandary. One might contend that undermining those arguments for rejecting (S)'s truth is not by itself sufficient grounds for affirming it. Still it could be an opening for a positive case if the context of utterance can shape the relevant class for evaluation: that is, if T-Contextualism advances the cause of a relevant

¹⁸Could one avoid this by taking the existence-entailment of those properties as equally minimalist? That would require a massive re-reading of Crane's entire earlier exposition.

¹⁹But see my (2004)

alternatives thesis. Because we have foregone assistance from the most prominent recent analyses that would explain our tendency to regard (S) as true (see note 14), more must be said about this appeal to T-Contextualism. I now turn to that task.

6

The passage from Austin, Maud ‘in’ a parked car, and the variations on ‘Al lives in Chicago’ all strongly indicate that truth or falsity hangs on a range of circumstances that can vary with the context under which they are produced. Corroboration from a different direction is provided by John Searle (1980). He states that when one utters “Bill cut the grass”, we take this as saying that he mowed it, not that he stabbed it with a knife. Similar shifts take place for cutting hair, cutting a cake, and so on. But Searle’s wording of this point warrants additional clarification. He writes,

The view I shall be espousing is that in general the meaning of a sentence only has application (it only, for example, determines a set of truth conditions) against a background of assumptions and practices that are not representable as a part of the meaning (1980: 221).²⁰

If the “background of assumptions and practices” determines a set of potential truthmakers, it undoubtedly falls within T-Contextualism. Although it might seem possible to take these as determiners of the truthbearer itself, Searle correctly rules out their semantic interpretation (*viz.*, as contributions to the sentence’s meaning) to block that avenue. In sum, the proliferation of the unambiguous cases reviewed earlier make it hard to resist assimilating cases of this ilk to the other examples for T-Contextualism. The shifting truth conditions for propositions fitting a schema like ‘x cuts y’ are, then, just an alternative way of expressing different sets of relevant contexts for a truthmaker.

But nothing thus far extends that view, as promised, to fictional utterances such as (S). I now turn to those.

For starters I agree with Crane that utterances relevantly like (S) are propositions and admit truth-values. Defending those claims would take me well beyond the scope of this essay, so for now I simply rest my case on their *prima facie* plausibility. My contention, then, is that on those assumptions T-Contextualism supports (S)’s capacity for truth *in the right context*.

On the surface a difference is needed to distinguish the status of (S) from that of Sherlock Holmes is a notorious lover (or, a barber, or a banker). It is obvious (S) has an aura of rightness about it that the alternatives lack. How otherwise could (S) be taken as something to be explained by virtually everyone studying this question, including those who don’t believe (S) expresses a proposition? Our job is to account for it. For starters (S)’s sort of rightness looks like the kind we ascribe to truths rather

²⁰Unlike the more modest version I am defending, Searle supposes this extends to *every* statement, not just to some of them. I need not take a stand pro or con on that extension.

than falsehoods. Of course appearances can be misleading for various reasons, but there must be reasons. One reason to withhold truth, or even its aptitude, might be the requirement that a referent exist, the clearly dominant view in the discipline, occasionally bolstered by a disparagement of contrary intuitions. However, in the opposite direction there is the pull of a school of thought that the understanding of *linguistic meaning* begins from what *we* mean, personal meaning; what our words mean must be understood as ultimately derivative, perhaps with restrictions, from what we mean.²¹ When combined with the irresistibility of the view that personal reference can be to certain things that don't exist, this tarnishes (if indeed it does not thoroughly undermine) the motivation for an existence requirement for reference in general. (More on that below.) But even if we regard a reference to Sherlock Holmes as *only* personal, once (S)'s status as a proposition is not put in question, there remains a strong proclivity to get things right under implied truth-conditions. How are we to do that?

A popular account of truth may be taken to overturn this appearance. Beginning with statements made with simple subject-predicate sentences it declares them true, as Quine has put it, "accordingly as the general term is true of the object to which the singular terms refer" (1960: 96).²² Crane calls this "the simple view of truth predication", stating that "[a] predication is true just in case the objects have the properties [predicated of them]" (2013: 21). Much hangs on the term 'object'. Undoubtedly, Quine would vehemently reject applying it to non-existents. But, concentrating on his formula rather than his intention, unless we have the backing of a particular philosophical theory, an ontologically loaded reading of the term's use is not a factor in its spontaneous universal acclamation. A commonplace understanding of the phrase need not be committal, and it would not be unnatural to say that the name in (S) refers to Sherlock Holmes without there being the slightest suggestion that Holmes is thereby more than a fictional character. For workaday conversation the simple view needn't imply the ontological status that many theorists find in it. In Crane's case this is emphatically so because he discusses nonexistent *objects* ("of thought", 46) with equanimity (2013: 16, see also my note 8). So it is far from certain that "the simple view of truth predication", read naively, imposes a condition on truth-evaluation incompatible with my earlier remarks.

Now T-Contextualism does not imply that (S) is true in every context. It would be false when the comparison is with

(I) Sir Ian Blair is a detective.

What I am claiming for it is that this is not the only context relevant to (S)'s truth-evaluation. No doubt it is a stretch to imagine a hearer who is familiar with the fictionality of Holmes but who lacks knowledge about Holmes being a consulting

²¹Grice (1957), Searle (1983), Recanati (1993).

²²The claim is regularly taken for a truism in philosophy. For one of many examples Jerrold Katz writes, "an atomic sentence is true if and only if the thing designated by its subject has the property designated by its predicate" (1966: 46).

detective. Nevertheless, for the statements of (S) that most naturally come to mind, and the one emphasized by Crane, one would need to muster up a less typical scenario to suppose Holmes is being confused with a real-life detective. There it is the predicate, not the subject, that is intended to carry its informational content. The matter in question seems not unlike a discussion that would take place in which two scholars disagree over whether certain claims about an Old Testament figure are supported by the text and historical background assumptions; but in which the figure's historical existence is not part of the conversation. Because the figure may have existed attaching a fictional preface to their claims would be wholly irrelevant to the differences between them. But this is not because they are committed to the figure having existed. Similarly, in the Holmes case the contrast of occupations is foregrounded in the situation in which we imagine (S) uttered. Once the existence of the subject is not part of what its assertion is intended to address, predicates come into play. Their informative content is acquired by belonging to a certain range of options, forming the context against which truth and falsity is contemplated. In fact, for T-Contextualism all one needs to claim is that this interpretation is possible, not that it is predominant. Its possibility renders it truth-eligible in the right context.

Accordingly, I find it difficult see the kind of 'getting it right' involved as anything other than the sort of thing that goes on in our truth evaluations. What makes it possible appears to be simply the context in which (S) is uttered. Here are a few ways, although not the only ones, to account for this.

- (i). Term reference's existence condition has exceptions; existence is a rule-of-thumb, a *prima facie* rather than absolute (/strict) requirement.
- (ii). Although reference has failed (less harshly, hasn't occurred), under exceptional circumstances, a proposition's truth will not be disqualified by not satisfying its requirement. The condition's failure will be outweighed by contextual factors.

What more can be said on behalf of each option?

Concerning (i), consider Crane's view that term reference is a technical notion: it imposes a certain restriction for systematic purposes not found in personal reference. It is unclear to me whether an independently devised theory of *truth*—that is, one not coming out of one's theory of reference with its built-in restrictions—is to be regimented along the very same lines. If one's truth theory contained restrictions that excluded intuitively compelling instances it too would be technical. However, even under that circumstance it is hard to see why truth's requirements should carve reality at the same joints as those peculiar to term reference. If we wish to understand our practice of dividing up propositions into truths and falsehoods—without being concerned exclusively with the use of the predicate 'is true', predicated in truths only on the rare occasion—what reason have we to believe that it should be circumscribed by the dictates of that theory of reference? No doubt referential restrictions of theory are intended to enhance understanding. But when the related but distinct worldly relation of truth is added to the mix such theories and understanding may diverge. In those cases we may forfeit a full comprehension of the way we acknowledge truth by adhering to the restrictions placed on a regimented referential system.

Truth's requirements seem to float independent of those for term reference: (i) nicely preserves this aspect of truth's independence.²³

Concerning (ii), stating, asserting, hypothesizing, and the like are speech-acts (even if only *sotto voce*), and their truth-evaluable contents are products of those acts. As early speech-act theorists noted, an act may succeed even when certain conditions for its optimal, perfect performance go awry. (Sincerity conditions are prime examples.) Similarly, when predicates have been foregrounded, we can expect the same to hold for truths in which, say, names of fictional characters are the failed referring expressions. One suggestion is that ideal cases of this kind of compensation exemplify the following two conditions:

Personal reference holds in cases in which term reference either fails or might fail, and in which those cases are grounded in something more than a mistaken theory.

(i) and (ii) have the virtues of preserving both the tenacious impression of truth and truth's independence from technical requirements designed specifically to regiment reference. This is even evidenced in the three strategies for fictional utterances I have rejected—*viz.*, fictional prefaces, fictional entities as abstract objects, and make-believe utterances. Each has tried to account for the persistent appearance of truth despite what they consider reference failure. This too adds to the pertinacious impression that views about truth do not march in lock-step with the requirements of term reference.

Resistance to (i) and (ii) may be grounded in a disputable, albeit standard, way in which the distinction between term and personal reference are portrayed by philosophers. I cannot do full justice to this issue here, but a brief summary may help us to understand why I believe questioning it is a good idea. Both forms have been taken from the start as well-understood, so that a natural next question for a further understanding each might be 'what are their appropriate linguistic slots'? Term reference is then regarded as falling under semantics and an existence requirement subsequently links it to an overwhelming array of truths of primary concern. Consequently, personal reference, not living up to that demand, is assigned to pragmatism or less. Call that view and its approach *taxonomy*.

But why start by taking term reference as fully-formed? Let us say that our two forms confront each other over a more fundamental difference, that of explaining the representational prowess of their devices. Following David Kaplan (1989) call it *metasemantics*. Personal reference traces that ability to the intentions of its users. No doubt such intentions must be qualified; an effort can fail. But even for materialists

²³I am aware that more must be said about the price to be paid by regarding the requirements for reference as less than absolute. For example, reference is generally taken to be a *relation*; accordingly, the relata of a referential vehicle must belong to a set-theoretic domain, for which existence is a standard condition of membership. Granting the first step, a comprehensive treatment would need to examine in much finer detail the requirements for domain membership. Perhaps existence is too facile—*relevance* might do for some subjects, or one or another form of free logic may be more suitable. However, I make no commitment on the issue, assuming that it should follow rather than precede the level of data collection at which the present discussion is pitched.

it is difficult to see how a primarily philosophical account could go much deeper. However, stand-alone term reference is not equally explanatory. Why are its tokens more than naked noises, like the sound of wind chimes? How are they capable of indicating (non-causally!) something beyond themselves? We need further details to understand an expression's representational prowess. Eventually it too must invoke users and their intentions at a more fundamental level, although, unlike personal reference, more in the way of further content conditions are needed. But this much is enough to rank term reference as derivative rather than original; its requirements enter not at the top, but via a more fundamental encumbrance. From a metasemantical perspective term reference's desiderata are best reflected not as demands of reference per se, but as conditions that logically-speaking precede it. This is in line with existence counting as no stricter than a defeasible requirement (as (i) predicted).

Even without this referential background no party to the present dispute is likely to deny the appearances I have been citing. But I want to emphasize that those who reject the force of such comparisons fail to realize just how deeply the evidence cuts despite its incompatibility with their preferred scheme. I can see no sensible option other than to allow that in the circumstances I have emphasized it is true that Sherlock Holmes is a detective. This may be accounted for by setting up a range of relevant alternatives in the predicate for determining truth.

What I can say about *relevant alternatives* here must remain sketchy. For example, there may be many such incompatible classes of alternatives within whatever boundaries are drawn. Perhaps the range is things that account for Holmes' celebrity, or, more simply, his livelihood, or perhaps things at which he excels, and the list goes on. Standard practice is to leave those choices open when further specification makes no difference on the occasion. Also, the relevance class might vary between speaker and hearer(s). Even where they extensively overlap, one participant's range may be narrower than another's. Nor do I foresee a way to obtain a principled rule for deciding between conflicting ranges. For colors the ranges are fairly clear; for occupations it is clear only that we can rule out those that would be anachronistic or too extraordinary for certain cultures (e.g. *fin de siècle* Great Britain had no television repairmen or tribal shamans). However, this is not a serious problem; not only because there will usually be considerable overlap between differing ranges, but also because any relevant alternatives account worth its salt should also reflect to a large degree cases of radical miscommunication and breakdown when they occur due to incompatible RAs. My general point is only that when discussing truth versus falsity, we make our judgments against the background of options that contribute directly to those verdicts. Not everything is up for grabs for every statement. Context matters. The special lesson of this exploration is that if there are propositions whose subjects are fictional, context also makes a difference for their evaluation.

References

- Austin, J. L. (1962). *How to do things with words*. London: Oxford University Press.
- Carston, R. (1988). Implicature, explicature, and truth-theoretic semantics. In R. Kempson (Ed.), *Mental representations: The interface between language and reality*. Cambridge: Cambridge University Press.
- Cartwright, R. (1987). Propositions. In *Philosophical essays*. Cambridge, MA/London: MIT Press.
- Crane, T. (2013). *The objects of thought*. Oxford: Oxford University Press.
- Crimmins, D., & Perry, J. (1989). The prince and the phone booth: Reporting puzzling beliefs. *The Journal of Philosophy*, 86, 685–711.
- Donnellan, K. (1966). Reference and definite descriptions. *The Philosophical Review*, 75, 281–304.
- Dretske, F. (1970). Epistemic operators. *Journal of Philosophy*, 67, 1007–1023.
- Dretske, F. (1972). Contrastive statements. *Philosophical Review*, 81, 411–430.
- Grice, P. (1957). Meaning. *The Philosophical Review*, 66, 377–388.
- Grice, P. (1989). Logic and conversation. In *Studies in the way of words*. Cambridge, MA: Harvard University Press.
- Harari, Y. (2014). *Sapiens: A brief history of humankind*. London: Harvill Secker.
- Kaplan, D. (1989). Afterthoughts. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan*. New York/Oxford: Oxford University Press.
- Katz, J. (1966). *The philosophy of language*. New York: Harper & Row.
- Kripke, S. (1977). Speaker's reference and semantic reference. *Midwest Studies in Philosophy*, 2, 255–276.
- Lewis, D. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, 8, 339–359.
- McGinn, C. (2000). *Logical properties*. Oxford: Oxford University Press.
- Nunberg, G. (1979). The non-uniqueness of semantic solutions: Polysemy. *Linguistics and Philosophy*, 3, 143–184.
- Quine, W. v. O. (1960). *Word and object*. Cambridge, MA: The MIT Press.
- Recanati, F. (1993). *Direct reference*. Oxford/Cambridge, MA: Basil Blackwell.
- Recanati, F. (2005). Literalism and contextualism: Some varieties. In G. Preyer & G. Peter (Eds.), *Contextualism in philosophy: Knowledge, meaning, and truth*. Oxford: Oxford University Press.
- Schiffer, S. (2003). *The things we mean*. New York: Oxford University Press.
- Searle, J. R. (1980). The background of meaning. In J. R. Searle, F. Kiefer, & M. Bierwisch (Eds.), *Speech act theory and pragmatics*. Dordrecht: D. Reidel Publishing.
- Searle, J. R. (1983). *Intentionality*. New York: Cambridge University Press.
- Sperber, D., & Wilson, D. (1985/1986). Loose talk. *Proceedings of the Aristotelian Society*, 86, 153–171.
- Sperber, D., & Wilson, D. (1986). *Relevance*. Oxford: Basil Blackwell.
- Urmson, J. O. (1976). Fiction. *American Philosophical Quarterly*, 13, 153–157.
- Van Inwagen, P. (2000). Quantification and fictional discourse. In A. Everett & T. Hofweber (Eds.), *Empty names, fiction and the puzzles of non-existence*. Stanford: CSLI Publications.
- Vision, G. (1993). Fictionalism and fictional utterance. *Pacific Philosophical Quarterly*, 43, 344–357.
- Vision, G. (2004). *Veritas*. Cambridge, MA: MIT Press.
- Walton, K. (1990). *Mimesis as make-believe*. Cambridge, MA: Harvard University Press.
- Zalta, E. N. (2000). Pretense theory and abstract object theory. In A. Everett & T. Hofweber (Eds.), *Empty names, fiction and the puzzles of non-existence*. Stanford: CSLI Publications.

Subsentential Speech Acts: A Situated Contextualist Account



Joanna Odrowąż-Sypniewska

Abstract Subsentential speech acts are subsentential utterances by means of which one “makes moves in a language game”, i.a. asserts, requests, asks, etc. I’m going to propose a pragmatics-oriented account of such acts based on Recanati’s moderate relativism. *Pace* Recanati and following Perry, I’ll argue that at least in the case of subsentential speech acts we have to postulate unarticulated constituents in explicit contents as well as in the situations of evaluation. I’ll suggest a test which helps to determine which constituents belong where. I’ll also argue – somewhat paradoxically – that subsentential speech acts are not an argument in favour of contextualism in any interesting sense.

Keywords Contextualism · Minimalism · Moderate relativism · Situated unenriched illocutions · Subsentential speech acts · Unarticulated constituents

1 Introduction

The most commonly given examples of subsentential speech acts are expressions such as “Nice dress”, “From Spain”, “Where?” used in such circumstances in which speakers uttering them are regarded as “making moves in a language game” (Dummett 1973: 194), e.g. asserting, asking, promising, etc. If Mum knows that Johnny is looking for his ball and says “Under the table” she probably means – and will be understood as meaning – “Your ball is under the table” (or “The ball is under the table”, “It is under the table”, etc.). If “Under the table” were uttered in response to the question “Where is my ball, Mummy?”, Mum’s utterance would be a straightforward case of ellipsis. However, since in our example it is uttered in a discourse initial position, with no (articulated) linguistic antecedent, regarding it as an ellipsis is much more problematic.

J. Odrowąż-Sypniewska (✉)
Institute of Philosophy, University of Warsaw, Warsaw, Poland
e-mail: j.odrowaz@uw.edu.pl

Philosophers who discuss such subsentential utterances might roughly be divided into (i) those who think that such utterances are not speech acts, because either their content is not determinate enough or they do not have determinate force; (ii) those who claim that they are speech acts but are not really subsentential, since they can be treated as a special kind of ellipsis (because either their antecedents are raised to salience by extralinguistic context (see Stanley 2000) or they belong to a limited group of ellipses which do not need antecedents (see Merchant's limited ellipsis hypothesis (2010))); and (iii) those who think that despite being genuinely subsentential they are full-blooded speech acts, because they have propositional and truth-evaluable contents (most notably Robert Stainton, but also e.g. Robert Elugardo, Ellen Barton, Alison Hall).¹ I side with this last, "pragmatics-oriented" (Stainton 2006b: 228) group, and I'm going to suggest a situated contextualist account of subsentential speech acts based on Recanati's moderate relativism (see Recanati 2007), which, in my view, is a promising way of developing Stainton's insights.

The best known pragmatics-oriented account is due to Stainton. He argues that speakers can assert with subsentential expressions and provides multiple examples illustrating this fact. The literal content of such acts is propositional even though what the subsentential utterance expresses is not:

what is asserted when a subsentence is used communicatively is that proposition which results from minimally adding to the content of the bare phrase actually uttered so as to arrive at a proposition (Stainton 2006b: 60).

According to Stainton in order to grasp the propositional content of a subsentential speech act the hearer needs first to grasp both the content of the utterance and the content from the extralinguistic context and next combine them "by function-argument application" (Stainton 2006b: 156). He distinguishes two cases: either the speaker utters an expression whose content is the argument while the context provides the propositional function for this argument or the speaker's utterance's content is the propositional function in which case the context has to provide an argument for this function. The content from the context is never translated into natural language, hence the resultant propositions are neo-Russellian in nature. The hearer knows that he must look for content in the extralinguistic context, since he assumes that the speaker intends to communicate a proposition but notices that what is uttered does not express any. The hearer recognizes that

¹In fact the picture is more complex: for instance, neither Stanley (2007) nor Corazza (see below) fit easily to these groups. Stanley suggests a "divide and conquer" strategy, according to which alleged subsentential speech acts belong to one of three categories: they are either ellipses, whose linguistic antecedents are raised to salience by context, or are not speech acts because their content and/or force is not determinate, or else are shorthands. See Stanley 2007. His view was thoroughly criticised by Elugardo and Stainton (2004). I also need to add that both Merchant and Stainton admit that some examples of subsentential utterances escape their proposed analyses. Thus, for instance, Stainton writes: "My burden is merely to argue that syntactic ellipsis is not the whole story; I'm happy for it to be part of the story" (2006a: 94), while Merchant observes that there are nonsentential structures that are not "amenable to an elliptical analysis" (2006: 73).

her interlocutor could mean such-and-such proposition—got from combining the language-derived content with a salient entity of some kind. And so on, in familiar Gricean fashion (Stainton 2006b: 158).

The problem with this explanation is that Grice hasn't been concerned with the way the hearer grasps *what is said*. As we know, Gricean what is said is to be meant by the speaker, truth-conditional and close to conventional meaning. The role of the context is very limited: it only helps to disambiguate and to supply reference for indexicals. The context and the inferences made by the hearer are much more important in deriving implicatures from what is said but such implicatures do not count as literal content of assertions. Stainton, however, wants the content that the hearer arrives at to be asserted – and not merely otherwise communicated. In my opinion Stainton is right that the content of subsentential utterances must be supplemented from the situation in which such utterance is taking place, but there is a better way of explaining how it happens: instead of appealing to Gricean mechanisms, one should turn to Perry's ideas concerning the completion of thoughts.

2 Situation-Relativity: Perry and Corazza

In his 1994 paper John Perry uses Wittgenstein's builder's language game to argue that nouns used in isolation (such as "Slab" or "Pillar") in the right circumstances might e.g. express the command (e.g. that the builder's assistant is to pass a slab) even though such words are not part of sentences. He claims that there are two ways of expressing complete thoughts (Perry 1994: 33): either one can utter a complete sentence whose elements correspond to all necessary constituents of a proposition or one might merely "complete a thought": i.e. utter a word or a phrase which supplies only the missing constituents necessary to express a thought whose other constituents are part of the situation of utterance. Using Perry's 1980 terminology we may say that those other elements present in the situation are unarticulated constituents and things that the subsentential utterance *concerns* rather than *is about*.²

More recently Eros Corazza (2011) has used Perry's framework to argue that subsentential speech acts are best seen as situated unenriched illocutions. He appeals to Perry's distinction between reflexive and incremental truth-conditions and to his idea of situations providing necessary elements and argues that we do not need to

²According to Perry there are also unarticulated constituents that are part of the content. These are the things that the proposition *is about*. So for instance, if I say "It's raining" while looking out of the window, the place I'm talking about (i.e. the place at which I make my utterance) is something that my utterance concerns, but if I say this thinking of a different place, then that place (which is cognitively represented by me) will be a thing that the utterance is about. See below.

appeal to ellipsis or pragmatic enrichment to claim that subsentential illocutions might be successful communicative acts.

Corazza considers the following scenario:

John, a well-known anti-Fregean, has been told that Jane is desperately looking for Dummett's *Frege: Philosophy of Language*. Jane walks into John's office. John suddenly utters:

(4) Hidden on top of the shelf.

Jane reaches to the top of the shelf and grabs Dummett's book. (Corazza 2011: 570)

He claims that the reflexive truth-conditions of the subsentential expression "Hidden on top of the shelf" uttered in the above situation are the following:

There is an x such that

- (i) (4) concerns x
- (ii) x is hidden on top of the shelf. (Corazza 2011: 575)

In his view it is enough for the communication to succeed if the hearer grasps a general thought expressed by a subsentential utterance (i.e. its reflexive truth-conditions).³ If John says to Jane "Hidden on top of the shelf", his communicative act will be successful if Jane entertains the thought "There is something hidden on top of the shelf". She doesn't have to grasp a singular thought expressed by John (such as "Dummett's book is hidden on top of the shelf"). Even if Jane forgot that she had been looking for Dummett's book, grasping the reflexive content of John's utterance will enable her to find the book on top of the shelf. John's utterance need not be pragmatically enriched, for it is "situated" (Corazza 2011: 577) and all the necessary elements are already present in the situation.

There are two main problems with Corazza's view interpreted as an account of subsentential speech acts. Firstly, Corazza is more interested in successful communicative acts than with speech acts. However, communication may be successful even if no words are uttered, whereas in order to talk about speech acts we need an utterance with a more or less determinate content and illocutionary force.⁴ By saying that such utterances need not be enriched, Corazza likens them to nudges and kicks under the table, which indeed might be communicatively successful but cannot be regarded as speech acts. We may imagine that instead of uttering "Hidden on top of the shelf" John merely points with his head to the relevant place. If his gesture were salient enough, it is quite probable that Jane would understand that he intends to convey that the book is on top of the shelf. In such a case John's action would constitute a successful communicative act, but it would not be a *speech* act. Plausibly illocutionary acts may be divided into linguistic and non-linguistic acts,

³*Nota bene* Corazza – following Perry – claims that one grasps reflexive truth-conditions just by being linguistically competent.

⁴I do not mean by this that the content has to be fully determinate. Thus, unlike Stanley (2007: 44), I'm prepared to grant that "a thirsty man who staggers up to a street vendor and utters: (...) water" has made a speech act.

but only the former can be regarded as speech acts. Moreover, even among linguistic illocutionary acts there will be those that do not qualify as full-blown speech acts. For instance, in order for a subsentential utterance to be an assertion it must have asserted as well as otherwise communicated content. And Corazza would be the first to acknowledge that “There is something hidden on top of the shelf” is not the asserted content of John’s utterance.⁵

The second problem is that the way in which the reflexive truth-conditions are generated is unclear. One may ask how do we know that the reflexive truth-conditions for subsentence “Hidden on top of the shelf” include “*there is an x such that (. . .) x is*”. For someone who regards such utterances as ellipsis the answer is easy, but we are not told how someone who does not think that subsententials are fragments of larger partly unpronounced structures arrives at such reflexive truth-conditions. The situation in which the exchange is taking place cannot do all the work, especially that Corazza allows the possibility that Jane might have forgotten that she is looking for Dummett’s book. A related problem is that Corazza considers in any detail only two subsentential utterances: “Hidden on top of the shelf” and “Reserved” (uttered by a waiter pointing at a table in a restaurant). In both cases it is relatively easy to supply reflexive truth conditions (“There is something hidden on top of the shelf” and “This table/It is reserved” respectively). However, there are many utterances for which it would be a challenge. For instance, John might have said only “On top of the shelf”. If we do not enrich his utterance and the situation is not determinate enough, it will be unclear whether its reflexive truth conditions are “There is something on top of the shelf” or, for instance, “Put something on top of the shelf”.

Moreover, it is not clear whether understanding only the reflexive truth-conditions of an utterance is a sufficient guarantee of the communicative act being successful. For Corazza an act is successful if as a result the hearer acts according to the speaker’s intention.⁶ Imagine, however, that Jane indeed forgot that she had been looking for Dummett’s book. If she grasps only the content “There is something on top of the shelf” then it is hard to predict what she will do. If she grasps the incremental truth conditions of the utterance, then she will realize that John is talking about the book for which she was looking and she will reach to get it. However, if the only thing that she understands is that something is on top of the shelf, then she may simply ignore John’s remark, in which case it will not be a successful communicative act.

⁵Corazza stresses that Perry’s pluri-propositionalism does not amount to the claim that more than one proposition is necessarily expressed by one sentence. Reflexive truth conditions are not the proposition expressed. It is incremental truth-conditions that play the role of asserted content. See Corazza 2011: 564.

⁶More precisely, Corazza writes that “a communicative act can be said to be successful inasmuch as the speaker and her audience are co-situated and the action-related output of the speaker and her audience is positive” (Corazza 2011: 578).

Thus, it seems to me that Corazza's view cannot be regarded as a successful attempt to analyse subsentential utterances as illocutionary acts, let alone as speech acts.

3 Moderate Relativism

In my opinion there is a better way of using Perry's ideas in accounting for subsentential speech acts. The way to do it is to adapt Recanati's framework of moderate relativism. Recanati's relativism presupposes two principles: duality and distribution, which say that to assign a truth-value to a proposition we need both a content and a circumstance of evaluation (duality), and that the determinants of truth-value distribute over content and circumstance (distribution). To those two principles Recanati adds the principle of economy, which says that the richer the content, the poorer the circumstances can be and vice versa. (Recanati 2007: 33–34) Thus, for instance, the content of "It's raining here and now" uttered on the 27th January 2017 in Paris will be that it is raining on 27.01.2017 in Paris and the circumstance will consist only of the possible world, whereas the content of "It's raining" uttered in the same situation will be that it is raining, and not only the world but also time and place will belong to the circumstances. The explicit content (which Recanati calls *lekton*) plus the situation of evaluation add up to the complete content (aka the Austinian proposition). Recanati claims that his framework allows to account for the conflicting intuitions we have concerning Barwise and Etchemendy's example of Claire the player (see Recanati 2007: 50). In the example the speaker is watching a game of poker. He thinks that Claire is among the players and says "Claire has a good hand now". As it turns out Claire is playing poker but in a different part of town. Moreover, she does have a good hand. Moderate relativism allows us to say that the sentence uttered by the speaker is true (since Claire indeed has a good hand at the time of utterance), but the utterance is not true (because it concerns the situation that the speaker observes and Claire is not present in that situation).

Recanati argues at length (2007: part 9) – *pace* Perry (1986) – that *lekta* do not contain any unarticulated constituents. According to him, all unarticulated constituents are things that the proposition concerns (in Perry's sense) and hence belong to the circumstances of evaluation. Perry argues that invariant, externally determined unarticulated parameters are constituents of the situation, whereas those that are not invariant and are cognitively discriminated are constituents of the content. Recanati claims that invariability can be so extended as to include also cases in which unarticulated constituents are cognitively determined, so there is no need to posit unarticulated constituents in the *lekton*.

4 Situated Contextualism

In my view Recanati's moderate relativism provides a good framework to analyse subsentential speech acts.⁷ In order to assign subsentential utterances truth-conditions we need both the explicit content and the circumstances of evaluation. The determinants of truth-value are distributed over the linguistic content and the circumstances, and since the linguistic content is so impoverished, the circumstances must be quite rich for such utterances to qualify as speech acts. The two-level evaluation that Recanati applies to the Claire example, can easily be adapted for subsententials: "From Spain" uttered by John holding a letter is not a sentence, hence it has no propositional content and no truth-value. Nevertheless, seen as an utterance it can be assessed as true or false in a particular discourse situation. However, it seems to me that in the case of such acts we have to postulate a two-staged principle of distribution: it is not only the case that the determinants of truth-value distribute over content and circumstance but also the content itself is distributed over the locutionary what is said and the situation of the utterance. In other words, I claim – *pace* Recanati and following Perry – that at least in the case of subsentential speech acts *lekta* have unarticulated as well as articulated constituents. Postulating only such unarticulated constituents that the proposition concerns will not do. I agree that we understand utterances "It's raining" and "It's three o'clock" without even thinking about the place and time of the utterance and about time zones respectively. The situation of the utterance provides all the relevant parameters and we do not need to consider them; they need not be cognitively articulated. Therefore, the proposition does not have to be *about* them and they may be 'left' in the situation. The case of subsentential utterances is different, however. It doesn't suffice that the missing elements are in the situation: they have to be cognitively articulated and have to be part of the content of the utterance. If John utters "From Spain" holding a letter, but his utterance is not accompanied in John's thoughts by any representation of the letter, then John cannot be said to have grasped the relevant proposition. If John doesn't have any representation of the object he's talking about, his utterance cannot be regarded as a speech act. And it seems to me that it is not enough that the letter is cognitively articulated in some other mental representation of the speaker. Regarding Perry's scenario in which his son in Palo Alto – after talking to his older son in Murdock on the phone – says "It's raining", Recanati notices that since Murdock is not an element of the situation of the utterance, it has to be articulated "in some mental representation in the speaker's mind" (Recanati 2007: 226). This mental representation need not be identical to the belief which is expressed by the utterance "It's raining", however. It may be some other mental representation in the speaker's mind. Thus, even though for the utterance "It's raining" to concern Murdock, Murdock has to be cognitively

⁷In *Perspectival Thought* (2007: 252) there is one example of subsentential speech ("Very handsome!" used as a comment on someone's appearance), but as far as I'm aware Recanati has never discussed subsentential speech acts in any detail.

articulated in some mental representation of the speaker, the utterance may concern Murdock (and not be about it), because the representation in which it is articulated need not be the one which is expressed by the utterance. As an example where a similar phenomenon occurs Recanati gives the following two sentences “Berkeley is a nice place. There are bookstores and coffee shops at every corner” and says that Berkeley is an articulated constituent of the content of the first statement, but an unarticulated constituent of the second (in his sense, i.e. Berkeley is a constituent of the situation and a thing that the second statement concerns in Perry’s terminology) (Recanati 2007: 226). In my view it is very doubtful whether the second statement merely concerns Berkeley. It seems to me that it is more plausible to regard it as a case of ellipsis with the locution “in Berkeley” omitted, but clearly determined by the linguistic context.⁸ Be that as it may, the case of subsentential utterances is different. We are assuming that they appear in discourse-initial positions, so without any linguistic antecedents in which the missing elements could be articulated. As for the mental representation, it is clearly not sufficient for a subsentential utterance to count as a speech act that the object the utterance supposedly is about is mentally represented somewhere in the speaker’s mind. Imagine that Mary comes home and sees John standing in front of the mirror trying on a new jacket and holding a letter. If John utters “From Spain” it will not be clear whether he means the letter or the jacket, so – I argue – his utterance will not be a speech act, for its content will not be determinate enough. In order to count as speech act it has to be either about the letter or about the jacket, so one of these objects needs to be an unarticulated constituent of the content of John’s utterance and a thing that his utterance is about and not merely concerns.

The account I advocate may be called ‘situated contextualism’, for on the one hand it postulates unarticulated constituents in the explicit content (and not only in the situation as relativism does), and on the other hand it stresses the importance of the situation that the utterance concerns: the elements of the situation are included both in the content and in the circumstances of evaluation.

It seems that it might be possible to use a disquotational indirect reports test⁹ to distinguish those unarticulated constituents that the proposition is about (i.e. those that belong to the *lekton*) from those that it concerns (i.e. elements of the situation). If an unarticulated constituent has to be added in order to make the indirect report correct, it should be regarded as a constituent of the *lekton*. Thus, since we may report someone’s utterance “It’s raining” just as “He said that it was raining”,

⁸I agree with the remark that Recanati makes in footnote 98 that Perry’s original description of his scenario also makes it seem like an ellipsis, for the son utters “It’s raining” in reply to Perry’s question: “How are things there?” (see Perry 1986). In order to exclude the ellipsis analysis we may modify the scenario slightly and assume that the son’s utterance is not made in reply to any question and has no linguistic antecedent.

⁹The inter-contextual disquotational indirect reports test has been proposed by Cappelen and Lepore (2005: 88) for an entirely different purpose: as one of their context-sensitivity tests (which supposedly demonstrates that expressions like “ready”, “tall” or “knows” are not context-sensitive). Nb. In my opinion this test does not deliver the results that Cappelen and Lepore wanted.

unarticulated constituents (place, possibly time) are not elements of the *lekton*; they are items that the proposition concerns. On the other hand, since in order to report John's utterance "From Spain" the reporter must add the missing object (e.g. "John said that *the letter/it* was from Spain"), the letter should be counted as an unarticulated constituent of the *lekton*.

5 Are Subsentential Speech Acts an Argument for Contextualism?

Subsentential speech acts have been regarded as an important argument in the debate between minimalism and contextualism (see e.g. Stainton 2006b, Merchant 2010). Stainton argues that "[t]he existence of genuine nonsentential speech acts, and nonsentential assertions in particular, entails that [contextualism is true]." (Stainton 2006b: 227). Stanley, who prefers a semantics-oriented approach – claims that genuine nonsentential speech acts – if there were any – would be the evidence that the truth-conditional role of context is not limited to disambiguation and resolution of indexicality, since it would provide constituents "directly to what is asserted" (Stanley 2000: 402).

Although I propose a contextual account of such acts and claim that there are unarticulated constituents in their content, I think – somewhat paradoxically – that they are not an argument in favour of contextualism in any interesting sense. There are two main ways of interpreting the debate between semantic minimalism and contextualism. One way is to regard it as a dispute concerning whether or not semantics delivers intuitive truth-conditions; another way is to see it as a discussion concerning whether or not semantics delivers propositional contents. If we understand it in the first way, i.e. the way, which takes the debate to be about what is asserted, then indeed subsentential speech acts demonstrate that contextualism is correct. This way of understanding the debate is clearly assumed by both Stainton and Stanley (as evidenced by the quotes given above).¹⁰ However, I take this understanding to be a mistaken one, since on this interpretation of the debate everyone is a contextualist save those who advocate indexicalism (see e.g. Stanley's *Introduction* to 2007).¹¹ In particular, current-day semantic minimalists (see Borg 2012) have given up on the claim that asserted content is not pragmatically enriched. Borg claims for instance that according to minimalism semantics should deliver "an account of the meaning of sentences rather than an account of our intuitive judgements of speech act content" (Borg 2012: 51).

¹⁰This understanding of the debate is also accepted by Recanati in *Truth-Conditional Pragmatics*, where he says that the only *substantial* debate between truth-conditional pragmatics and minimalism is the one concerning the determinants of the intuitive what is said (Recanati 2010: 13–14).

¹¹Compare here Recanati's remarks on I-Minimalism (2010: 14; esp. footnote 11).

If on the other hand we interpret the debate in the second way, as concerning semantically determined truth-conditional what is said – as I think we should interpret it – then nonsentential speech acts, even if genuine, have no direct bearing on its outcome. By definition, they are not sentences, so they have no truth-conditional content whether semantically determined or not. What is asserted in them needs to be pragmatically enriched, but this doesn't mean that what is said is also so enriched. Thus, if we agree with Merchant (2010: 152) that “Stainton's real goal here [i.e. in *Words and Thoughts* – JOS] is to show that at least ‘moderate’ contextualism is correct: (...) this is the claim that context (and pragmatics) determines at least part of what is *said* (or ‘sentence meaning’) in addition to what is *meant* (or ‘speaker meaning’)”, then we must conclude that Stainton's attempt has failed. One may argue that although subsentential speech acts demonstrate that speaker meaning (what is asserted) is partly determined by context, they do not show that what is said is in this way determined.¹² Hence, such a view is compatible with minimalism as well as contextualism. In order to make subsentential speech acts an argument for contextualism one would have to argue that what is said is propositional as well. However, such an argument has yet to be put forward.

Acknowledgements The research for this paper was possible thanks to the support of National Science Centre, Poland (grant OPUS 8; research project UMO-2014/15/B/HS1/00171).

References

- Borg, E. (2012). *Pursuing meaning*. Oxford: Oxford University Press.
- Cappelen, H., & Lepore, E. (2005). *Insensitive semantics*. Oxford: Blackwell Publishing.
- Corazza, E. (2011). Unenriched substential illocutions. *Philosophy and Phenomenological Research*, 83, 560–580.
- Dummett, M. (1973). *Frege: Philosophy of language*. London: Duckworth
- Elugardo, R., & Stainton, R. (2004). Shorthand, syntactic ellipsis, and the pragmatic determinants of what is said. *Mind & Language*, 19(4), 442–471.
- Merchant, J. (2006). “Small structures”. A sententialist perspective. In L. Progovac, K. Paesani, E. Casielles, & E. Barton (Eds.), *The syntax of nonsententials: Multidisciplinary perspectives* (pp. 73–91). Amsterdam/Philadelphia: John Benjamins.
- Merchant, J. (2010). Three types of ellipsis. In F. Recanati, I. Stojanovic, & N. Villanueva (Eds.), *Context-dependence, perspective, and relativity* (pp. 141–192). Berlin: The Gruyter Mouton.
- Perry, J. (1986). Thought without representation. *Proceedings of the Aristotelian Society*, 60, 137–151.
- Perry, J. (1994). Davidson's sentences and Wittgenstein's builders. *Proceedings and Addresses of the American Philosophical Association*, 68(2), 23–37.

¹²Stainton argues that nonsentential assertions have propositional literal asserted content but he does not claim that phrases used to make such assertions express propositions. On the contrary, he says that e.g. “Moving pretty fast” “doesn't express a proposition, *even after reference is assigned to indexicals and such*”, “*even once its contextual parameters have been saturated*, it still refers to a function from objects to truth-values” (Stainton 2006b: 17, his emphasis). For his understanding of the debate and of the notion of what is said see Stainton 2006b: 223–225.

- Recanati, F. (2007). *Perspectival thought. A plea for (moderate) relativism*. Oxford: Oxford University Press.
- Recanati, F. (2010). *Truth-conditional pragmatics*. Oxford: Oxford University Press.
- Stanley, J. (2000). Context and logical form. *Linguistics and Philosophy*, 23(4), 391–434. Reprinted in Stanley 2007.
- Stanley, J. (2007). *Language in context*. Oxford: Clarendon Press.
- Stainton, R. (2006a). Neither fragments nor ellipsis. In L. Progovac, K. Paesani, E. Casielles, & E. Barton (Eds.), *The syntax of nonsententials: Multidisciplinary perspectives* (pp. 93–115). Amsterdam/Philadelphia: John Benjamins.
- Stainton, R. (2006b). *Words and thoughts: Subsentes, ellipsis, and the philosophy of language*. Oxford: Oxford University Press.

Some Constraints on Contextualism About Modals



Daniel Skibra

Abstract The following paradigm, most closely associated with the work of Angelika Kratzer, dominates the literature on the semantics of modals: **MUST** is a universal quantifier over a contextually determined set of worlds, where $\lceil \text{MUST } \phi \rceil$ is true just in case ϕ holds in each of these worlds. The quantifier is restricted via the setting of values for various contextual parameters relative to which the modal is interpreted. Notwithstanding its important virtues, I aim to show that the standard way of carrying out the paradigm overgenerates in a pretty severe way with prejacent with eventive predicates. The argument of this chapter concerns this overgeneration problem and can be broken down into two sub-claims: **CLAIM 1**: If context is responsible for setting the values of the parameters relative to which modals are semantically evaluated, then we can't account for this common pattern in the interpretation of modal auxiliaries. **CLAIM 2**: These data provide independent support for recent revisionist approaches in linguistics, and places constraints on the way we think of context as influencing the interpretation of natural language modals.

Keywords Modals · Event semantics · Aspect · Epistemic modals · Deontic modals · Context · Contextualism

Thanks to Tadeusz Ciecierski and Paweł Grabarczyk for editing the present volume. This material was presented at the first Context, Communication, and Cognition Conference in Warsaw in June of 2016. I thank Michael Glanzberg, Adrian Briciu, Cem Şişkolar, and Zsofia Zvolensky for helpful comments on my talk.

D. Skibra (✉)

Department of Philosophy, Northwestern University, Evanston, IL, USA

1 Introduction

The following paradigm dominates the literature on the semantics of modals: **MUST** is a universal quantifier over a contextually determined set of worlds, where $\lceil \text{MUST } \phi \rceil$ is true just in case ϕ holds in each of these worlds. The quantifier is restricted via the setting of values for various contextual parameters relative to which the modal is interpreted. This paradigm originates in the work of Angelika Kratzer¹ and it has a number of important virtues. It allows for an elegant and general picture of modals according to which modals form a unified linguistic system. According to this picture, modals are not systematically ambiguous, but have a uniform kernel of meaning which they contribute to the sentences that contain them.

Notwithstanding these virtues, I aim to show that the standard way of carrying out the paradigm in philosophy and in much of semantics gets a simple yet crucial detail systematically wrong. The detail concerns patterns of modal interpretation. On this “Standard View”, the meaning the modal comes to have in a sentence owes ultimately to contextual factors. So, we’d expect contextual variation to provide us with certain readings. But in fact, we do not see this; certain modal interpretations, which the standard view predicts, are simply unattested. The Standard View faces an overgeneration problem. The argument of this paper concerns this overgeneration problem and will be broken down into two sub-claims:

CLAIM 1: If context is responsible for setting the values of the parameters relative to which modals are semantically evaluated, then we can’t account for a common pattern in the interpretation of modal auxiliaries.

CLAIM 2: These data provide independent support for recent revisionist approaches in linguistics, and places constraints on the way we think of context as influencing the interpretation of natural language modals.

2 The Semantics of the Standard View

Modals in natural language encompass a variety of auxiliary verbs like *must*, *might*, *may*, *ought*, *can*, *should*, etc., adjectives (*possible*), adverbs (*possibly*), some main verbs.² My concern here is mainly with the auxiliaries – in particular, with the way in which modal auxiliaries come to exemplify the different classes of meanings they can have. It is standard to call these classes “flavors” of modality. Many modal auxiliaries admit of multiple flavors. Take *must* as it occurs in the different sentences in (1).

¹Cf. especially Kratzer (1977, 1981, 1991, 2012).

²An example of main verbs exhibiting modality would be drives in *This car drives 150 mph*.

- (1) a. (In view of our evidence) Sam must be the murderer.
 b. (In view of the laws) You must not steal other people's property.
 c. (In view of your desires) You must try the chocolate cake.
 d. (In view of your goals) You must avoid the Dan Ryan Expressway at rush hour.

In all the sentences in (1), *must* has a different flavor. We would say that in (1a) the flavor is epistemic, in (1b) it is deontic, and in (1c) and (1d) it is bouletic and teleological, respectively. The free-relative *in view of...* phrases make the flavor of modality particularly prominent. Let's call this ability of a modal to exhibit different flavors of modality *flavor diversity*. Within each flavor, there are further distinctions in meaning modals can come to exhibit. For example, holding fast an epistemic flavor, (1a) may differ in meaning depending on whose information state is at issue; if it is the knowledge available to some group, for example, or to the speaker. Let's call this *perspective diversity*.

In a moment, I will discuss the putative sources of both flavor diversity and perspective diversity. Let's first rehearse the semantics for modal sentences as given by Kratzer. A modal sentence is analyzed as having the following logical form: $\ulcorner \text{Mod}(R)(\phi) \urcorner$. The modal takes two arguments; *R*, the restrictor, which determines the domain the modal quantifies over, and ϕ , the nuclear scope. In the literature, it is also common to refer to the nuclear scope as the *prejacent*, and roughly speaking, it corresponds to the unmodalized sentence, which the modal scopes over like an operator. Insofar as modals are like quantifiers at some level of abstraction, the restrictor restricts the quantifier domain. It can be explicitly expressed; in fact, Kratzer suggests that the *in view of...* phrase serves to express the restrictor. More commonly, however, the restrictor is not explicitly expressed. It remains an argument of the modal, but is supplied contextually.

The restrictor is itself comprised of two elements that Kratzer calls the *conversational backgrounds*. In fact, on Kratzer's developed theory,³ there are two conversational backgrounds making up the restrictor, the modal base (*f*) and the ordering source (*g*). Both of these are functions from possible worlds to sets of propositions (the nature of which I will describe momentarily). Since propositions are themselves characterized as sets of possible worlds, this amounts to saying that *f* and *g* are functions from possible worlds to sets of sets of possible worlds. Ultimately, the modal base contributes to picking out an initial domain of possible worlds, which the ordering source restricts further by imposing a partial order on the value of the modal base.

This division of labor between the modal base and the ordering source makes for another benefit of the Standard View; it can account for the distinction, long noted in generative linguistics, between epistemic and root modals. Epistemic modals correspond to those having an epistemic flavor in the flavor taxonomy I

³Cf. Kratzer (1981).

offered above. By contrast, the other flavors of modality mooted above (deontic, bouletic, teleological, etc.) are all root modals. Kratzer's semantics neatly captures the distinction between epistemic and root modals by saying that they simply have a different kind of modal base. The former have a modal base that characterizes a particular information state. An epistemic modal base picks out a set of propositions (each of which is a set of possible worlds) characterizing the information state. Root modals have a circumstantial modal base; this modal base picks out a set of propositions (again, each of which is a set of possible worlds) that characterize a particular circumstance. The following analogy would be apt: epistemic modals : epistemic modal bases :: root modals : circumstantial modal bases. Taking the generalized intersection of these sets of propositions gives us a single set – the set of worlds consistent with the circumstances, or the information. So these modal bases can give us the initial domain the modal quantifies over as follows.

- (2) a. $\bigcap f_{epistemic}(w) = \{w' \mid w' \text{ is compatible with what is known in } w\}$
 b. $\bigcap f_{circumstantial}(w) = \{w' \mid w' \text{ is compatible with certain circumstances in } w\}$

Though we won't be worrying much about the ordering source, let me explain how it works. The ordering source, g , also picks out a set of propositions, and \leq_g imposes a preorder on the worlds in $\bigcap f$ according to the propositions of $g(w)$ that hold at each world. g might pick out a set of propositions that characterize the laws, a person's desires, her goals, etc. We can think of g as further restricting the domain as follows: $BEST_g$ is a function that picks out the best worlds (according to g 's standards) in the domain.⁴

Putting the ingredients together, we get the following semantics for $\lceil \text{MUST } \phi \rceil$.

- (3) $\llbracket \text{MUST } \phi \rrbracket^{c,w,f,g} = \{w \mid BEST_{g(w)} \cap f(w) \subseteq \llbracket \phi \rrbracket^{w,f,g}\}$.⁵

Alternatively, we can put this in terms of truth-conditions as follows.

- (4) $\llbracket \text{MUST } \phi \rrbracket^{c,w,f,g} = \forall w' \in BEST_{g(w)} \cap f(w): \phi(w') = 1$

Reflecting on the roles of the conversational backgrounds, we see that flavor is often jointly determined by the modal base and ordering source. However, the roles of f and g in determining modal flavor aren't entirely symmetrical. An epistemic modal base is sufficient to make a modal have an epistemic flavor. Modals with a circumstantial modal base are roots, but which of the root flavors will depend on the kind of ordering source they receive. For example, if g picks out a set of propositions that characterizes a person's desires, then the modal will be bouletic, and so on.

⁴Though this diverges from Kratzer's presentation somewhat, it is now a fairly standard way of characterizing the ordering source. Cf. Portner (2009) for such a presentation.

⁵Note: I've suppressed other parameters of the interpretation function besides those relevant to our discussion.

Another thing we see is that these conversational backgrounds are contextually determined. What determines whether a particular modal base (for example) is circumstantial or epistemic is purely contextual, limited only by the metasemantics of context sensitivity. It is to say that flavor and perspective diversity owes entirely to context sensitivity. This results in a fairly strong claim; absent some extenuating circumstances, context alone is responsible for determining the flavor of a given occurrence of a modal. Though this has long been thought to be a feature of the Standard View as opposed to a liability, it is this strong claim about the role of context that will cause trouble for the Standard View.

3 Noteworthy Features of the Standard View

I'd now like to note a few features of the Standard View. Some of these were discussed in passing in the previous sections, but they deserve to be emphasized. These features aren't so much entailments of the account as they are independent theoretical commitments of the view. First of all, there is the view that modals are not ambiguous. That is to say, the phenomenon of flavor diversity does not warrant positing different but homophonous modal words; a deontic *must*, and an epistemic *must*, as the case may be. Part of the problem with an ambiguity approach to flavor diversity is that perspective diversity seems to pose the same kind of challenge. Even within a particular flavor, a modal sentence-type could admit of different interpretations on different occasions of use. In *Sam must be the murderer*, holding fast an epistemic flavor, this sentence could mean that, Sam's being the murderer follows from my evidence, or that his being the murderer follows from our joint evidence, and so on. The point, made in Kratzer (1977), is that if the mere perceived difference in meaning between (*inter alia*) the deontic and epistemic interpretations of a modal can justify positing ambiguities, then parity of reasoning suggests we ought to follow the same methodology in accounting for the second kind of meaning difference. But this multiplies the different senses of *must* intolerably. To continue with the example of epistemic modals, it seems like we would have to posit a different *must* for each possible information state the modal can be assessed against. Much better is to say that modal words are not ambiguous at all – they have a uniform kernel of meaning. There is one *must*, one *may*, etc. This commitment is captured by something like UNIFORMITY.

UNIFORMITY: Modal words in natural language have uniform lexical entries, in spite of their ability to express different types of modality.

But if modals are not ambiguous, then flavor diversity must owe to something else besides the kind of lexical ambiguity described above. Kratzer's semantics gives us a way to characterize this uniform kernel of meaning, and suggests a natural way to explain flavor (and perspective) diversity. One of the selling points of Kratzer's semantics is that it allows us to make good on this theoretical commitment. There is

a natural way to characterize the uniform kernel of meaning that modals have; they have a lexical entry that is composed of various parameters. A modal contributes this lexical entry to the logical form of the sentence that contains it, which we can characterize as in (5). Assigning values to these parameters is what winds up determining the different meanings of the modals.

$$(5) \quad \llbracket \text{must} \rrbracket^{w,f,g} = \lambda P \lambda f \lambda g \lambda w [\forall w' \in \text{BEST}_{g(w)}(\bigcap f(w)) : P(w') = 1]$$

The next theoretical commitment to be discussed started coming into relief with our discussion of the implementation of UNIFORMITY. It suggests that the ultimate meaning of the modal is determined purely by the contextual provision of values to these parameters. That is to say, the modal having the uniform kernel of meaning it apparently has, resulting differences in meaning are then purely a matter of context sensitivity.

PRAGMATIC RESOLUTION: The contextual parameters introduced by modals exhaustively determine the meaning of the modal along the flavor dimension through the pragmatic assignment of values to the parameters.

I suggested earlier that neither of these features, UNIFORMITY or PRAGMATIC RESOLUTION, were entailments of the view – that instead they were independent theoretical commitments of the view. Nothing in the Standard View requires the kind of strong claim about context-sensitivity embodied by PRAGMATIC RESOLUTION, but given that f and g are contextual parameters, it is difficult to see any other alternative. If values for f and g are contextually supplied, and if assigning values to f and g determine modal flavor (and perspective), then it seems as though flavor diversity is simply a matter of context-sensitivity, pure and simple.

4 Putting Pressure on Pragmatic Resolution

We are now able to describe in more detail the problem for the Standard View I aim to show. Given the semantics of the Standard View, and its two theoretical commitments, it would predict that a sentences like (6) and (7) can have epistemic as well as root (say, deontic or teleological) interpretations.

(6) John must be at the store.

(7) John must go to the store.

Indeed, we can give (6) an epistemic or a deontic reading as the context might require. For example, imagine a context wherein we arrive at John's apartment and wonder where he is. We note that his shopping cart is gone and utter (6). This makes for a natural epistemic reading. For a deontic reading, imagine a context wherein we note that John has an obligation to buy groceries for his household (suppose it's his turn to do this), and we utter (6), noting that the store is where he needs to discharge

this obligation. There is little difference between (6) and (7), so we'd expect (7) to similarly be amenable to epistemic and root interpretations as well. Curiously, however, this is not what we see. I will sharpen this observation momentarily, but for now I'd just like to note that getting an epistemic interpretation for (7) is very difficult, unlike for (6), for which there was no difficulty at all in accessing an epistemic reading.

The first thing to note about how (6) and (7) differ is that the predicate of the prejacent in (7), *go to the store*, is eventive.⁶ In (6) it is stative. Without delving into the literature on lexical aspect too much, we can draw an intuitive contrast that will make the present point clear. Some predicates describe events and some describe states.⁷ Eventive predicates, like *run*, *go to the store*, *draw a circle*, *eat an apple*, etc., are all dynamic; they happen, and their happening takes time. Stative predicates, like *love (a person)*, *think*, *have a cold*, *own a car* or *hear birds*, by contrast, either hold or do not hold; they do not *happen*. Insofar as they do hold, they lack eventive's dynamism, and hold homogeneously over an interval. The reason for drawing this contrast will become apparent shortly.

According to the Standard View, there may be lexical constraints on the kinds of flavors a modal can express. For example, it is often pointed out that *might* has no deontic reading. This is thought to be due to *might*'s lexical idiosyncracies. In terms of the Standard View's framework, *might* simply doesn't accept the right kind of conversational background to generate a deontic meaning; $[[\text{MIGHT } \phi]]^{c,w,f,g}$ is not defined for deontic g . This is simply a quirk of *might*; it is part of its uniform kernel of meaning that it doesn't take a g characterizing deontic flavor, which ensures that *might* never has this reading.⁸ Consider the modal *must*, as in (7). It can express epistemic as well as a range of root meanings. Unlike *might*, it is not lexically prevented from expressing deontic or epistemic flavors of modality; (6)'s admitting of both epistemic and deontic interpretations shows that this is so. Absent such lexical constraints, the Standard View doesn't give any kind of story about how we might effectuate a restriction on what kind of conversational background a modal

⁶That is, non-stative. Cf. Vendler (1957) and Dowty (1979), for the classic categories of lexical aspect/*Aktionsart*. Cf. Kenny (1963) and Ryle (1949) for important precursors.

⁷I'm following the distinction made by Vendler (1957) and Dowty (1979) who categorize processes as kinds of events. Other authors have differed on this, making processes their own metaphysical category intermediate between states and events. (Cf. e.g., Mourelatos (1978). But cf. Smith (1999) for an argument that processes ("activities" in the Vendler/Dowty taxonomy) are types of events.) Since the contrast that seems most relevant to the data will be that between stative and non-stative predicates, whether activities/processes are events or not isn't terribly important. What is important is that they are not states.

⁸Perhaps it isn't even defined for circumstantial f . This is a stronger claim, and its plausibility would depend on other factors as well, like whether the "metaphysical" reading of *might* diagnosed by Condoravdi (2002) is a root modal or not. Condoravdi considers it a non-root modal, but assigns it a modal base much like the circumstantial one. This would complicate the epistemic/non-root : epistemic modal base :: root : circumstantial analogy, but I will not address the matter further here, and flag it only to set it aside.

can take. So, we are left with the prediction that (7) ought to admit of epistemic as well as root interpretations.

However, as noted above, it is very difficult to get an epistemic reading of (7). Root readings of (7) come without much difficulty, as it is easy to envision natural scenarios wherein an utterance of (7) might have a deontic or a teleological reading.⁹ (7) is simply typically not uttered to convey an epistemic modal sentence. With some effort, one could dig up an epistemic reading, but there are constraints on how this might be done, and these constraints are significant. At this point, it should become clearer that the difficulty is not specifically with (7) but with all sentences where the prejacent has an eventive predicate. If we replace the predicate in (7) with another eventive predicate, like those given earlier as examples, the pattern persists; it remains difficult to get an epistemic reading of *must* whenever the predicate is eventive.

We can now say more about the constraints under which *must* in (7) (and sentences like it) can have an epistemic reading. We've been considering *episodic* readings of these eventive sentences; such sentences report specific events (even if the specific events are non-actual, as they are in modal sentences). Epistemic readings are only available on **habitual** and **futurate** readings of the predicate. They vanish on episodic (and non-futurate) readings of the predicate. A first way to get an epistemic reading of (7) is on a non-episodic interpretation of the predicate, but on a **habitual** reading. Habitual sentences do not report particular events; rather, they report regularities of events.¹⁰ With a habitual interpretation of the predicate, we may get a corresponding epistemic reading of *must*, as in (8).

(8) John must go to the store (Mondays/often/every day after work).

Note that the habitual reading of (8) is made more prominent when it is accompanied by a frequency adverbial, such as *Mondays*, *often*, or *every day after work*. However, the frequency adverbial is optional. It might be harder to access the habitual reading without it, but a context where a person's habits are under discussion would suffice to make the habitual reading prominent even without the adverbial.¹¹

Another option for an epistemic reading would be with a so-called **futurate** interpretation of the predicate. This interpretation is a little more delicate, and (7) isn't the best example to illustrate such a reading. Futurates are sentences that refer to some future event or state of affairs without any overt future-

⁹Deontic: John and his roommates divide household responsibilities and take turns going to the store. The cupboards are bare and it's John's turn to go to the store. His roommate utters (7). Teleological: John is making a cake and his roommate is helping him by collecting some ingredients from the cupboards. He notices that John lacks milk and butter for the cake. He utters (7).

¹⁰Cf. Carlson (2005) or Krifka et al. (1995) for an overview on habituals.

¹¹Consider a discussion where we're considering what it is that John does Mondays, and note that his cupboards are always well-stocked that evening. The habitual reading of *John must go to the store* would be prominent even without any frequency adverbial.

referring morphology on the verb. That is, they refer to the future without any of the standard markers for doing so.¹² For example, *The plane leaves at 5 o'clock* can refer to a future time, specified by *at 5 o'clock*, without having to use the stereotypical way English refers to future events (i.e., with the auxiliary *will*). The contrast between this sentence and (7) also serves to make clear the difficulty of getting this reading with a predicate like *go to the store*. In order for the utterance of a futurate to be felicitous, the event described needs to be planned or pre-determined to happen in advance. It is often said that they have a “scheduled” reading.¹³ A flight’s leaving is indeed typically scheduled, making the futurate reading of this sentence easy to access. Events of going to the store are much less often scheduled or planned in the same way, and a futurate reading of a sentence containing this predicate is commensurately difficult to access. A context where John’s store-going is understood to be planned in advance could, in theory, yield a futurate reading of the predicate. However, it is much easier to access such a reading if the predicate stereotypically lends itself to a planned interpretation. Since leavings are very often planned (especially if they involve catching a flight or a train), it is easy to access a futurate interpretation of a sentence like *John leaves tomorrow*. Futurate sentences can allow for an epistemic interpretation of *must* when the embedded predicate is eventive as in (9).¹⁴

(9) (Context: John’s suitcases are by the door) John must leave tomorrow.

To summarize to lesson of this section: epistemic readings of (7) are available when the prejacent has a futurate or a habitual interpretation, but not on episodic, non-futurate interpretations of the predicate. As acknowledged before, strenuous effort might be required to think of a context that would support such an interpretation, but notwithstanding the effort, such readings exist. Aside from these two kinds of interpretations of the predicate, however, we simply do not see epistemic readings of a modal like *must* when the predicate is eventive.

Although the predicates themselves (*go to the store*, *leave*, etc.) are lexically eventive, habituals and futurates have a special kind of semantics. The semantics for these constructions is fairly complex and I won’t discuss the details much, for the point I intend to make doesn’t depend on many of these details. The important points here are two-fold. First, the most prominent accounts of the semantics of these

¹²An early discussion of futurates occurs in Lakoff (1971). Cf. also Vetter (1973) and Huddleston (1977) for discussion. The most comprehensive analysis to date is due to Copley; cf. especially her (2008, 2009), and subsequent work.

¹³Cf. Kaufmann (2002, 2005) for more on this “scheduled” reading.

¹⁴In light of the difficulty of accessing a futurate interpretation with *go to the store*, I’ll simply change the predicate in the example as opposed to thinking up a context that will support the futurate interpretation of the prejacent in (7). But it’s important to see that in principle, nothing prevents (7) from admitting of a futurate reading if the context allowed for it.

constructions would have it that they involve some kind of operator or element in the logical form which trigger these readings of the predicates to begin with; the variable binding operator GEN in the case of habituais, and the a phonologically null modal or futurate operator in the case of futurates.¹⁵ Secondly, both habituais and futurates in fact have a stative semantics.¹⁶ What these points jointly suggest is that even though the most deeply embedded predicates are eventive as a matter of their lexical aspect, the intervention of the habitual or futurate operators renders them stative. This observation is interesting for our present purposes because it is precisely the eventive predicate that seemed to make it so hard to for *must* to have an epistemic reading. But even when we were able to dig up epistemic readings for (7) (or a sentence relevantly enough like it), the condition under which these sentences allowed epistemic readings of *must* was that the prejacent could no longer be eventive.

We are then faced with a putative generalization; ***must with a prejacent containing a bare eventive predicate cannot have an epistemic interpretation.***¹⁷ I'll call this the "eventivity constraint." From the perspective of the Standard View, this is truly puzzling. After all, it is the *conversational background* that should determine the flavor of the modal. Since the conversational background is determined contextually, it should be a matter of context whether *must* is, say, epistemic or deontic in (7). But careful attention to these patterns of interpretation suggest it is not in fact a matter of context – at least, not the way context is generally presumed to determine the reading of the modal. (7) (and sentences like it that have eventive predicates) can only have an epistemic reading if the eventive predicates have readings that require an intervening element that stativizes the resulting expression. Context might serve to disambiguate, say, habitual from non-habitual prejacents, or futurate from non-futurate prejacents, but once the prejacent is determined to be episodic and non-futurate, context cannot conspire to produce an epistemic reading of *must*.¹⁸ In terms of the framework outlined in previous sections, this means that when the prejacent is episodic/non-futurate, *f* cannot have a value as an *epistemic* modal base.

¹⁵Cf. Krifka et al. (1995) for discussion of habituais and Copley (2009) and Kaufmann (2005) for futurates.

¹⁶Cf. Krifka et al. (1995) on this point for habituais, and Copley (2014) for futurates.

¹⁷By "bare", I simply mean to suggest that no obvious futurate or habitual operator intervenes.

¹⁸Another way of putting the point, using the terminology of Stanley (2000), is that when we interpret (7) epistemically, context is not playing its typical *truth-conditional* role in determining modal flavor, which is what it does when it provides values to the contextual parameters. Rather, context can only play this role after it first plays a *grammatical* role in disambiguating, say, habitual from non-habitual prejacents. Only if the prejacent is stative can the modal be interpreted with an epistemic flavor, so way in which the epistemic reading of (7) is obtained is not solely through context playing a truth-conditional role.

5 Accommodating the Data with a Selectional Restriction

Though she is not addressing the data I'm concerned with here, Kratzer (1981) raises the possibility of the lexical entries for modals coming with selectional restrictions – that is, with restrictions on the kind of arguments they can take. Accordingly, one may think that selectional restrictions give one an easy way to explain the eventivity constraint. In this section, I will argue that the best explanation of the eventivity constraint ought not rely on selectional restrictions.

An early account in generative linguistics for how lexical items are inserted into syntactic structures had it that items appear in the mental lexicon as a structured bundle of features, and that some lexical items put restrictions on which sorts of words they can take as arguments, over and above their subcategorization requirements.¹⁹ For a simple example, consider (10).

(10) John drank a brick.

Since *drink* is a verb, according to the grammar, it should receive a Noun Phrase argument, and *a brick* is in fact a Noun Phrase.²⁰ But there's something strange about (10) nonetheless; a brick isn't a thing you can drink! (10) is syntactically well-formed, at least insofar as its verb takes a Noun Phrase complement, and verbs subcategorize for this kind of argument. Nonetheless, it is unacceptable.²¹ The idea here is that the verb *drink* takes an argument that must have something like [+liquid] as one of its lexical features. *Water* has this feature (among others, like [–inanimate]), and so the lexical item *water* can be inserted into the N slot of the NP. You could extend the idea to *must* by saying that epistemic *must* puts a restriction on its arguments. In this case, epistemic *must* would be restricted to taking complements that have a [–eventive] predicate. However, only epistemic *must* would have this restriction. Deontic *must* would not have it, because it can have a [+eventive] complement, and in respecting UNIFORMITY, we are assuming that there is only one *must*.

If the appeal to a selectional restriction follows the traditional model I've sketched above, where a word or lexical item selects another kind of item based on the presence of absence of a certain feature, then this appeal to a selectional restriction for *must* is just a capitulation of UNIFORMITY. While the contribution

¹⁹Cf. Chomsky (1965) for an early explanation of the role of selectional restrictions. Cf. Freidin (1992) for a clear, text-book introduction.

²⁰I haven't said anything about what constitutes a grammar, and won't currently take a detour through such a discussion. For the sake of the example, we can presume something like a phrase structure grammar which includes the rule $VP \rightarrow V + NP$.

²¹Chomsky's famous sentence *Curious green ideas sleep furiously* is a relevant example. Part of the import of this sentence was that, though it is perfectly syntactically well-formed, it is purportedly uninterpretable. Selectional restrictions aim to do some work in explaining why this is. Note that selectional restrictions are not overridden on figurative interpretations of sentences. The theory would have it, rather, that in figurative interpretations of sentences, lexical items are required to interpret certain sentences like (10), in order to respect the selectional restriction.

must makes to the logical form is the same for root as for epistemic *must*, one flavor of *must* has a feature the other lacks. The advocate of a selectional restriction is smuggling in a difference-maker surreptitiously.²² Unlike *might*, whose idiosyncrasy does not result in multiple lexical entries, this proposal would split the lexical entry for *must* into one entry for *must_{epistemic}*, and one for *must_{root}*, with the former having a unique selectional restriction as one of its features. Perhaps one could avoid this contravention of UNIFORMITY by positing a relational selectional restriction. The relational selectional restriction for *must* would presumably be something like RELATIONAL RESTRICTION.

RELATIONAL RESTRICTION: if a modal sentence S has the form MOD(R) ϕ , where ϕ is the nuclear scope of S, MOD is *must*, and ϕ is [+eventive], then f in R is *f_{circumstantial}*.

Since RELATIONAL RESTRICTION applies uniformly to *must* it conforms to UNIFORMITY. Nonetheless, as an explanation of the patterns of modal interpretation captured by the eventivity constraint, invoking RELATIONAL RESTRICTION is deeply unsatisfying. For one thing, it simply restates the data. We also see no independent motivation for it. As a *place-holder* for a more robust explanation, I have no objection to RELATIONAL RESTRICTION. If, on the other hand, we simply stipulate that RELATIONAL RESTRICTION is part of the lexical meaning of *must* in order to explain the eventivity constraint, the stipulation is ad hoc.

There is a further, methodological reason to be wary of an explanation by RELATIONAL RESTRICTION. Selectional restrictions were originally underwritten by a theory of lexical insertion which explained how lexical items are inserted into the syntactic structures they figure in, but this theory of lexical insertion is no longer in favor. So, it had better be that positing a restriction of this sort is independently motivated, since selectional restrictions themselves are no longer an integral part of some piece of grammatical theory. Otherwise it is utterly mysterious what just what this restriction is.

This completes my case for CLAIM 1. I submit that from the perspective of the Standard View, there is an epistemic reading “missing” for (7) and sentences like it. The Standard View predicts an epistemic reading for (7) where none is attested, and lacks an obvious way to constrain the value of the modal base without

²²There is another worry, which I’ll mention just to put aside. The features I allude to in my discussion of (10) concern features *lexical items* have as part of their entries in the mental lexicon, if you will. The kind of feature one would need to appeal to leverage selectional restrictions in explaining the eventivity constraint are of a different sort. They would be features of *complex expressions* or clauses. The advocate of this response owes us a story about where these features come from, and if they come from the lexical items themselves, how complex expressions come to inherit them from the lexical items that compose them. The theory of lexical insertion I appealed to in my explanation of selectional restrictions is outdated and long superseded. More contemporary versions of Phi Theory or analyses dealing with agreement phenomena may be able to provide a story about the distribution of such features in the clause. However, as far as I am aware, none of these extant theories appeal to a feature like [+eventive], and the introduction of such a feature for the purpose of explaining the eventivity constraint would be ad hoc.

contravening UNIFORMITY. This is not an isolated, idiosyncratic kind of sentence, but is in fact a pretty common pattern in modal interpretation, as it involves not just (7), but all prejacentes with eventive predicates. The Standard View seems unable to accommodate this pattern of interpretation, predict it, or to explain it.

6 The Event-Relative Approach to Modal Semantics

I'll now make the case for CLAIM 2. This is the claim that the eventivity constraint provides independent support for a revisionist approach in the semantics of modals. The one I have in mind is that of Valentine Hacquard.²³ This is because adopting this revisionist approach allows us to accommodate most of the paradigm that motivates the Standard View, and makes explaining the eventivity constraint more tractable. As we will see, it respects UNIFORMITY, but it circumscribes the role of context relative to the strong commitment to context sensitivity encapsulated by PRAGMATIC RESOLUTION.

Contemporary approaches to semantics are likely to involve a neo-Davidsonian element, incorporating event variables into logical form. Recent work in linguistics has revised the standard Kratzerian semantics on precisely this front. On Hacquard (2010)'s approach, modals are analyzed according to the Kratzerian paradigm, except that modals take an *event* argument as opposed to a *world* argument.²⁴ Following the work of Cinque (1999) on the relative positions of functional heads, it is largely thought that the relative syntactic positions of modals is as follows (irrelevant projections omitted).

MODAL_{epistemic} > TENSE > ASPECT > MODAL_{root} > VP.

This means that a difference in flavor correlates with a difference in syntactic position. Note that from the perspective of the unmodified Standard View, this is itself a curious fact. One of Hacquard's crucial amendments to the Standard View is to re-define the conversational backgrounds with event arguments. Instead of the definitions for modal bases in (2), we have (11).

- (11) a. $\bigcap f_{epistemic}(e) = \{w' \mid w' \text{ is compatible with } \text{CON}(e)\}$ ²⁵
 b. $\bigcap f_{circumstantial}(e) = \{w' \mid w' \text{ is compatible with certain circumstances of } e\}$

²³Cf. especially Hacquard (2006, 2010).

²⁴Hacquard's analysis was meant to explain the apparent difference in syntactic position between root and epistemic modals and the so-called actuality entailments (Cf. Bhatt 1999) exhibited in certain languages. I'll merely be presenting the first half of these motivations, since this is most relevant to the current problem.

²⁵CON is a function from events to their contents, which is only defined for events that "have content" like illocutionary events or intentional events. Cf. Pietroski (2000) for more on CON.

The move allows one to derive the different flavors from the event argument the modals take. The syntactic positions of modals are constrained by where they can get this event argument. Davidson (1967) famously suggests that verbs are to be interpreted as predicates of events. On this analysis, the verb *stab* in the sentence *Brutus stabbed Caesar with a knife* is interpreted as a predicate that has an argument position for an event variable, as in (12).²⁶

(12) $\exists e$ [Stabbing(e, b, c) & With(e , a knife)]

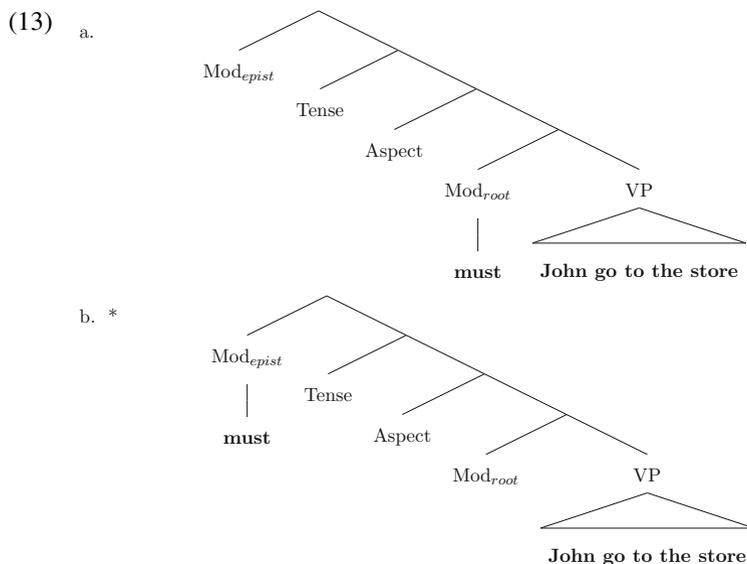
The reason “low” modals sit in the position just above the VP is that the modal base (and ordering source) take this event variable as an argument and yield circumstantial modal bases. They are therefore compatible with the range of root interpretations. Epistemic modals are “high” modals that sit above tense. They take an event variable as an argument as well, but not the event variable introduced by the VP. Instead, they are anchored to the *speech* event (the idea here, following Krifka (2001), is that high in the structure, there is an assertion operator, so the utterance-type contributes an event pronoun to the logical form), and such modal base projections yield an epistemic modal base.

Hacquard’s analysis doesn’t quite explain the eventivity constraint on its own; there is no immediately obvious reason why *must* cannot have an epistemic reading when the VP is comprised of an eventive predicate. But what is relevant to the present inquiry is the fact that the modal base parameter’s taking different arguments (either the event variable introduced by the VP or that introduced by the speech event) is what determines whether the value of the modal base is epistemic or circumstantial.²⁷ This allows a sub-sentential feature to influence the modal interpretation, albeit in an indirect way. We no longer have to ask why an eventive predicate would constrain the modal base parameter in some way. We can ask instead, knowing that epistemic modals sit in a “high” position and root modals in a “low” position, if there is some structural feature of these two configurations that would explain why there is a “missing” epistemic reading from the perspective of the Standard View. Recasting the puzzle in terms of Hacquard’s account of modal semantics, we need an explanation of why it is that (7) can have the syntactic

²⁶Davidson thought that the subject and object of the verb are themselves arguments of the verbal predicate along with the event variable. More contemporary versions of a Davidson-style event semantics separate these out via so-called theta roles. Cf. Parsons (1990) or Higginbotham (1985). This neo-Davidson approach to verbal arguments would render the sentence as follows: $\exists e$ [Stabbing(e) & Agent(e, b) & Theme(e, c) & With(e , a knife)].

²⁷In her dissertation, Hacquard distinguishes between “true deontics” and other roots. True deontics are supposed to sit “high” in the clause like epistemics, take a speech act event as an argument, and have a performative interpretation of the kind diagnosed by Ninan (2005). The existence of true deontics would be troublesome for me because true deontics seem subject to the eventivity constraint but would complicate the explanation I am about to give. For reasons I won’t go into here, I think this analysis of true deontics is a mistake, and in any event, it isn’t part of Hacquard’s proposal that has gained wider acceptance. The proposal about so-called true deontics doesn’t find its way into the (2010) paper, so I won’t complicate the explanation of the eventivity constraint to accommodate them, and will treat them instead as regular root modals.

representation in (13a) but not (13b). Why should an eventive VP prevent *must* from occupying the “high” position?



7 Towards an Explanation

Hacquard’s account allows us to describe the structural configuration of the various readings in (7) provided we make note of additional data. First, it is important to realize the role of grammatical aspect in (7). In a nutshell, grammatical aspect turns on the perfective/imperfective contrast, and refers to the presentation of the structure of the eventuality denoted by the predicate.²⁸ *Perfective* sentences describe *complete* eventualities, without regard for their internal structure. Conceptually, perfective aspect packages the event denoted by the predicate as a bounded whole, whereas imperfective aspect presents the event as in some way incomplete or ongoing. Imperfective sentences describe the eventualities from the inside, as it were, allowing language to account for the internal structure of the event. For example, in a sentence like *While I was writing a letter, Esther walked in*, the dependent clause is marked for imperfective aspect by way of the progressive, allowing for an interpretation whereby the event (in the main clause) of Esther’s walking in took place within the event of my writing a letter. That is to say, the writing of the letter was ongoing or incomplete at the point at which the walking in

²⁸Cf. Comrie (1976) for the classic textbook discussion.

event occurred. In English, the perfective aspect occurs in the unmarked case, where no imperfect marker (like the progressive, or habituality, etc.) is present. So a more complete representation of the structure of (7) would show that there is perfective aspect on the VP.

Episodic interpretations of eventive sentences in the present tense are unacceptable or anomalous. More specifically, given the discussion above, *perfective* eventive sentences in the present tense are unacceptable. There have been many proposed explanations for why this is. I won't survey the different proposals because the precise details don't quite matter for our analysis.²⁹ They do typically share a common thread, which has a rather intuitive motivation. The sentences in (14) have the common form $\lceil \text{PRES}(\text{PFV}(\phi)) \rceil$, where PRES is the present tense operator, and PFV is perfective aspect.

- (14) a. John goes to the store.
 b. Mike eats an apple.
 c. Mary pushes a cart up the hill.
 d. Tim wins the race.

Tense can be thought of as locating the time denoted by the predicate in relation to the time at which the utterance is made. What PRES does is indicate that the time of the predicate overlaps with the time of utterance.³⁰ The unacceptability of present tense (perfective) eventives is due to the fact that eventive predicates, which denote eventualities that take time to occur, can only be true at intervals, since the perfective ensures that the eventuality described is a completed one. By contrast, it is often assumed, going back to Taylor (1977), that the utterance time is either an instant or an interval so constrained that the event denoted by an eventive predicate cannot be contained in it.³¹ Eventive sentences take time to occur, so the events the sentences describe *couldn't* be true at an instant – these events take longer than an instant to occur. Stative predicates have what Bennett and Partee (1978) call the subinterval property, which means that for every interval at which the predicate

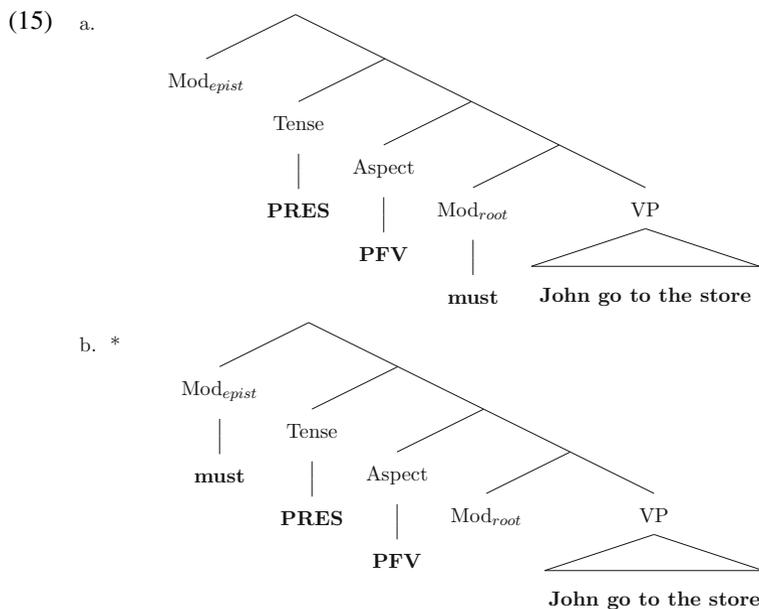
²⁹Some representative examples are Ogihara (2007), Parsons (1990), Cowper (1998), and Hallman (2009).

³⁰This is the traditional conception of tense. (Cf. Comrie (1985) for discussion.) A popular, alternative conception is due to Klein (1994), which departs from the traditional conception in important ways, but not *vis à vis* the underlying intuition I wish to illustrate here.

³¹Hallman (2009) holds the view that the 'now' of the utterance time is an instant, Bach (1981) and Parsons (1990) are examples of the view where it is a constrained interval – very small but not necessarily a singleton interval. Cowper (1998) argues that the utterance time is an instant, but that perfective aspect "packages" the eventuality as a point-like instant as well. She goes on to argue that the truth conditions of perfective eventives would require two instant-like points to coincide, which is so unlikely that languages avoid expressing this. Even the analyses that do not make the assumption that the utterance time is conceptualized as an instant tend towards an explanation that is conceptually similar to the one I am about to offer.

holds, it also holds at every subinterval contained therein. Stative predicates have no trouble being true under PRES: they can hold at the utterance time no matter the duration of the utterance interval. The sentences in (14) are acceptable when their predicates are interpreted habitually because habitual predicates are stative. The habitual interpretation of the sentences in (14) have the form $\lceil \text{PRES}[\text{HAB}(\phi)] \rceil$. PRES does not have an eventive complement since $\text{HAB}(\phi)$ is a derived stative predicate. Similar comments apply, *mutatis mutandis* to futurate sentences.

How does this help explain the missing epistemic reading in (7)? Adopting Hacquard’s account, the question that came into relief was why (7) cannot have the structure in (13b). I propose the following strategy for explaining the eventivity constraint, which makes essential use of the generalization about present tense (perfective) eventives: episodically interpreted eventive predicates are unacceptable in the present tense. When eventive predicates are not habituals or futurates, they are (i) perfective events that are (ii) directly exposed to PRES. (13a) and (13b) would be filled out as (15a) and (15b), respectively.



The explanation comes down to this. In order for *must* to sit in the high position (that of epistemic modals), it needs a stativizing aspectual operator to intervene (which both habitual and futurate operators can do, if they sit in the Aspect slot). If there is no such operator, then PFV will appear in the aspect slot, and we have a perfective eventive preajacent immediately in the scope of the PRES as in (15b), which is ruled out for the same reason that unmodalized present tense eventives are ruled out. Since (15a) is the only acceptable structure of the two, (7) has to receive a root interpretation of *must* when the predicate is interpreted episodically.

In a way, I've only offered impressionistic remarks of what a solution would look like by tying the problem to another one – the problem of why present tense eventive perfective sentences are bad. However, the point is that we already know we needed an account of what is wrong with present tense perfectives with eventive predicates. In fact, it doesn't quite matter what account we give of this; it could be imported seamlessly into the present framework to account for the distributional differences for modals. Interestingly, taking this approach means that there's a way in which the lack of epistemic modals with eventive prejacent does not have anything to do with the modals at all; these readings are ruled out on grounds orthogonal to the properties of modals, as long as there is a way to differentiate the syntactic position of modals as suggested by Hacquard's work. But in order to avail ourselves of this solution, we need to adopt the revisionist approach to modal semantics proposed by Hacquard (2010). This is the basis of my CLAIM 2.³²

8 Taking Stock

The present chapter argues that the Standard View overgenerates in a systematic way and offers the eventivity constraint as a generalization about the pattern of modal interpretation the Standard View fails to capture. Insofar as it takes *must* as its paradigm example, it raises the obvious question as to whether other modals exhibit the eventivity constraint or whether *must* is the only such modal. In fact, the question can be sharpened in a way that poses a genuine puzzle. For it seems that although (7) (repeated here as (16a)) precludes epistemic readings of *must*, (16b) admits of perfectly good readings of epistemic *may/might*, though the embedded predicates are identical.

- (16) a. John must go to the store.
 b. John may/might go to the store.

If the logical form is as I've said it is, and if my explanation for the lack of eventive-perfective epistemic readings of certain modals is correct, then my account leaves

³²Though the eventivity constraint is described in Steedman (1977), and though several important works in semantics discuss the sensitivity of modal interpretation to verbal aspect (Cf. Condoravdi (2002) for an example), the eventivity constraint has not received a lot of attention in semantics, to say nothing of philosophy. An exception is the work of Gillian Ramchand, which I became familiar with after offering the analysis outlined in this chapter. Ramchand (2014) makes a similar empirical observation and offers an analysis that has now expanded into the more recently published Ramchand (2018). Ramchand's account differs in a few ways from the one I sketch here; most obviously, she proposes a more radical revision of the Standard View. However, it agrees with the present proposal in a few important respects. Most importantly, it appeals to the different syntactic positions of root and epistemic modals to explain what I have called the eventivity constraint (what Ramchand calls "Epistemic-Stative Sensitivity"). A full engagement with Ramchand's work will have to wait for another time, but for now I will just flag the fact that her account corroborates some of the claims made in this chapter.

it mysterious why epistemic *might* and *may* seem perfectly acceptable with an eventive prejacent. Shouldn't they also be ruled out?

Though I will leave a detailed discussion of the difference between (16a) and (16b) to future work, I'd like to flag a possible explanation here. Notice that both of these *may* and *might* are possibility modals, and further, that even on an epistemic reading they have an obligatory future orientation. This much they share with the necessity modal *must*. Now, it *could* simply be that *must* exhibits the eventivity constraint while *might* and *may* do not, but then we'd need some additional explanation of why the combination of present tense and perfective aspect on an eventive predicate embed under a possibility modal when it was precisely the present tense + eventive-perfective configuration that ruled out the epistemic reading for *must*.

Here is where the future orientation of the complement clause becomes salient. In certain cases, we saw that a futurate expression could embed under epistemic *must* (recall (9), *John must leave tomorrow*). Examples like (9) actually conformed to the generalization I proposed because of the presence of a null futurate modal in the logical form, which has a stativizing effect on the complement. However, futurate readings of simple present tense sentences are fairly hard to come by, due to felicity constraints on the utterance of a simple futurate. It is often said that the predicate of a simple futurate needs to be "planned" or "scheduled" for the futurate to be felicitous. And this additional constraint explains why we simply don't encounter many sentences with a plausible epistemic-*must* + futurate configuration: there's no plausible planned reading of *go to the store* in (16a).³³

Here's the proposal. It may be that examples like (16b) are in fact just like (9) with respect to their logical form; they contain an embedded futurate which then stativizes the complement. This is why they have an obligatory future orientation. Where they differ from each other is with respect to the felicity conditions imposed by the futurate modal. When embedded under an epistemic necessity modal like *must*, the predicate *go to the store* has to have a plausible planned or scheduled reading in order for the futurate reading to be felicitous, just like in the case without the modal.³⁴ When embedded under an epistemic possibility modal, the prohibition against unplanned readings of the embedded predicate is relaxed. Of course, these remarks are highly speculative, largely because I haven't offered any explanation of why the prohibition is relaxed. But if something like this is on the right track, it offers a possibility for extending the account I present in this chapter for *must* to all modals. But spelling this out in satisfactory detail is a task I leave to future work.

³³In fact, with strenuous effort, you might be able to envision a context with such a planned reading. In such a context, you may see an epistemic reading of (16a).

³⁴This would predict that *John should/ought to go to the store* lacks an epistemic reading except on the planned/scheduled reading of *go to the store* as well. This sounds right to me, but my intuitions here aren't terribly strong.

To conclude, in this chapter I suggested a strategy for explaining the eventivity constraint that stays close to the paradigm underwriting the Standard View. My proposal has the following virtues. First, it preserves UNIFORMITY; *must* has a single lexical entry along the lines suggested by Hacquard.³⁵ Second, the epistemic reading of (7) is ruled out on independently motivated grounds. This account does not conform to PRAGMATIC RESOLUTION, which was shown to be too strong in any event. This, in turn, yields a surprising result. Context winds up not being the sole determinant of modal flavor. Rather, the grammatical environment of the event variable can determine part of the modal flavor (at least, it distinguishes epistemic from root flavors), leaving context to determine the remaining ingredients on the flavor-dimension via the ordering source. Insofar as the intuition motivating PRAGMATIC RESOLUTION concerned the context sensitivity of modals, the present account preserves what is right about this intuition while dispensing with what is wrong. Of course, the explanation sketched in this chapter is just the beginning of an account, but it shows the need for paying attention to the interaction of modals, tense, and aspect, and it serves to recapitulate Hacquard's point that context is not solely responsible for flavor-diversity – grammatical features play a role in ways underappreciated by philosophers following the Standard View and its competitors.

References

- Bach, E. (1981). On time, tense, and aspect: A essay in English metaphysics. In P. Cole (Ed.), *Radical pragmatics*. London: Academic.
- Bennett, M., & Partee, B. (1978). *Toward the logic of tense and aspect in English*. MS. Bloomington: Indiana University Linguistics Club.
- Bhatt, R. (1999). *Covert modality in non-finite contexts*. Ph.D. thesis, University of Pennsylvania.
- Carlson, G. N. (2005). Generics, habituals and iteratives. In A. Barber (Ed.), *Encyclopedia of language and linguistics*. Elsevier: Amsterdam, Netherlands.
- Chomsky, N. (1965). *Aspects of the theory of syntax*. Cambridge: MIT Press.
- Cinque, G. (1999). *Adverbs and functional heads: A cross-linguistic perspective*. New York: Oxford University Press.
- Comrie, B. (1976). *Aspect*. Cambridge: Cambridge University Press.
- Comrie, B. (1985). *Tense*. Cambridge: Cambridge University Press.
- Condoravdi, C. (2002). Temporal interpretation of modals. In D. Beaver, S. Kaufmann, B. Clark, & L. Casillas (Eds.), *The construction of meaning* (pp. 59–88). Stanford: CSLI Publications.
- Copley, B. (2008). The plan's the thing: Deconstructing futurate meanings. *Linguistic Inquiry*, 39(2), 261–274.
- Copley, B. (2009). *The semantics of the future*. New York/London: Routledge.
- Copley, B. (2014). Causal chains for futurates. In M. Kissine, P. de Brabanter, & S. Sharifzadeh (Eds.), *Future times, future tenses* (pp. 72–86). Oxford/New York: Oxford University Press.
- Cowper, E. (1998). The simple present tense in English: A unified treatment. *Studia Linguistica*, 52(1), 1–18.

³⁵And depending on how my speculative remarks at the end fare upon further investigation, the account may be perfectly general for modals.

- Davidson, D. (1967). The logical form of action sentences. In N. Rescher (Ed.), *The logic of decision and action*. Pittsburgh: University of Pittsburgh Press.
- Dowty, D. (1979). *Word meaning and Montague grammar: The semantics of verbs and times in generative semantics and in Montague's PTQ*. Studies in linguistics and philosophy (Vol. 7). Dordrecht: Springer.
- Freidin, R. (1992). *Foundations of generative syntax*. Cambridge: MIT Press.
- Hacquard, V. (2006). *Aspects of modality*. Ph.D. thesis, MIT.
- Hacquard, V. (2010). On the event relativity of modal auxiliaries. *Natural Language Semantics*, 18, 79–114.
- Hallman, P. (2009). Instants and intervals in the event/state distinction. MS. UCLA.
- Higginbotham, J. (1985). On semantics. *Linguistic Inquiry*, 16, 547–593.
- Huddleston, R. (1977). The futurate construction. *Linguistic Inquiry*, 8(4), 730–736.
- Kaufmann, S. (2002). The presumption of settledness. In *Proceedings of CLS* (Vol. 38, pp. 313–328).
- Kaufmann, S. (2005). Conditional truth and future reference. *Journal of Semantics*, 22(3), 231–280.
- Kenny, A. (1963). *Action, emotion and will*. New York: Humanities Press.
- Klein, W. (1994). *Time in language*. London/New York: Routledge.
- Kratzer, A. (1977). What 'must' and 'can' must and can mean. *Linguistics and Philosophy*, 1(1), 337–355.
- Kratzer, A. (1981). The notional category of modality. In H.-J. Eikmeyer & H. Rieser (Eds.), *Words, worlds, and contexts. New approaches in word semantics* (pp. 38–74). Berlin: De Gruyter.
- Kratzer, A. (1991). Modality. In A. von Stechow & D. Wunderlich (Eds.), *Semantik: Ein internationales handbuch zeitgenoessischer forschung* (pp. 639–650). Berlin: De Gruyter.
- Kratzer, A. (2012). *Modals and conditionals* (Oxford studies in theoretical linguistics). Oxford/New York: Oxford University Press.
- Krifka, M. (2001). Quantifying into question acts. *Natural Language Semantics*, 9(1), 1–40.
- Krifka, M., Pelletier, F. J., Carlson, G. N., Ter Meulen, A., Chierchia, G., & Link, G. (1995). Genericity: An introduction. In *The generic book*. Chicago: University of Chicago Press.
- Lakoff, G. (1971). Presupposition and relative well-formedness. *Semantics: An interdisciplinary reader in philosophy, linguistics, and psychology* (pp. 329–340). Cambridge: Cambridge University Press.
- Mourelatos, A. (1978). Events, processes and states. *Linguistics and Philosophy*, 2, 415–434.
- Ninan, D. (2005). Two puzzles about deontic necessity. In J. Gajewski, V. Hacquard, B. Nickel, & S. Yalcin (Eds.), *New work on modality* (Vol. 51, pp. 149–178). Cambridge: MIT Working Papers in Linguistics.
- Ogihara, T. (2007). Tense and aspect in truth-conditional semantics. *Lingua*, 117(2), 392–418.
- Parsons, T. (1990). *Events in the semantics of English: A study in subatomic semantics*. Current studies in linguistics (Vol. 19). Cambridge: MIT Press.
- Pietroski, P. M. (2000). On explaining that. *Journal of Philosophy*, 97(12), 655.
- Portner, P. (2009). *Modality*. Oxford: Oxford University Press.
- Ramchand, G. (2014). Stativity and present tense epistemics. *Proceedings of Salt*, 24, 102–121.
- Ramchand, G. (2018). *Situations and syntactic structures: Rethinking auxiliaries and order in English*. Cambridge: MIT Press.
- Ryle, G. (1949). *The concept of mind*. London: Hutchinson and Co.
- Smith, C. S. (1999). Activities: States or events? *Linguistics and Philosophy*, 22(5), 479–508.
- Stanley, J. (2000). Context and logical form. *Linguistics and Philosophy*, 23(4), 391–434.
- Steedman, M. J. (1977). Verbs, time, and modality. *Cognitive Science*, 1(2), 216–234.
- Taylor, B. (1977). Tense and continuity. *Linguistics and Philosophy*, 1(2), 199–220.
- Vendler, Z. (1957). Verbs and times. *The Philosophical Review*, 66(2), 143–160.
- Vetter, D. C. (1973). Someone solves this problem tomorrow. *Linguistic Inquiry*, 4, 104–108.

Index

A

Ambiguities, 3, 19, 108, 161, 171, 210, 247, 299
Answering machine puzzle, 7, 98, 99, 108, 110
Anti-intentionalism, 6, 82–88
Aspects, 2, 4, 5, 37, 44, 45, 47, 49, 75, 77, 81, 82, 100, 101, 135, 183, 187, 190, 195, 196, 201, 211, 232, 253, 280, 301, 304, 309–314

B

Background knowl-
edge/beliefs/assumptions/presuppositions,
1, 3, 4, 6, 8–10, 17, 18, 24–28, 31,
33–35, 37, 42, 47, 56, 59, 60, 62–64,
66, 67, 71, 72, 74, 78, 82, 86, 87, 105,
121, 126, 136–141, 146, 153, 156, 158,
159, 162, 164–174, 176, 183, 193, 194,
207, 211, 216, 218, 221, 222, 229, 230,
232–234, 237, 242, 248, 249, 251, 252,
254, 255, 257, 258, 260, 261, 266, 274,
277, 279, 289, 310

Belief

de re, 8, 26, 150, 157, 159, 162, 163,
169–171, 175, 176, 255, 258
de se, 8, 149–177

Binding, 111, 114, 122, 123, 149–177, 197

C

Characters, 1, 9, 16–18, 23, 24, 26, 27, 81, 97,
99, 108–110, 115, 127, 149–151, 153,
154, 156, 165, 166, 192, 196, 197, 269,
270, 272, 274–276, 278, 280

Common ground, 41, 56, 60, 64–67, 70, 74–76,
78, 249

Compositionality, 9, 10, 165, 166, 177, 179,
180, 184, 193, 194, 205–223

Content

mental content, 247, 250
propositional content, 138, 207, 208, 220,
221, 284, 289, 291

Context (cont.)

and circumstances of evaluation, 11, 127,
195, 288–290
and compositionality, 9, 179, 180, 184,
193, 194, 201, 208, 209, 212
context of utterance, 2, 11, 65, 66, 99,
108–110, 149, 152, 184, 207, 211, 213,
215–217, 219–233, 267, 276
context set, 3, 4, 19
contextual parameter, 2–4, 6, 7, 9, 30, 32,
34, 58, 67, 99, 206, 207, 212, 217, 292,
296, 300, 304

Contextualism, 9–12, 228, 229, 231–240, 249,
267, 289–292, 295–314

Coordination, 6, 40, 55, 83, 99, 126, 253, 256

D

Demonstratives, 6–8, 16–18, 21, 25, 29–31,
33, 38–43, 56, 58, 69, 73, 81–95, 99,
105, 113–115, 117, 125–146, 150, 154,
160, 166, 170, 177, 219, 220, 222,
265, 268

Demonstratum, 85, 93, 94, 115, 118, 219, 222

Direct reference, 9, 126, 150, 159–163,
180–187

E

- Episodes, 5, 16–19, 23, 25–27, 266
- Events, 7, 12, 25, 26, 44, 100, 106, 182, 229, 246, 247, 249, 252, 302, 303, 307–310, 314
- Event time (ET), 7, 100–104, 107–109, 111
- Existence, 6, 9, 129, 146, 149, 180, 182, 183, 193, 196, 199, 202, 206, 218, 223, 228, 266, 269–271, 274, 275, 278–281, 291, 308

G

- Gradable adjectives, 30–32, 36, 40, 44–46, 56–59, 68–71, 73–76

I

- Implicatures, 4, 10, 213, 258, 285
- Indexicals, 2, 15, 29, 59, 97, 116, 125, 149, 179, 207, 229, 265, 285
- Intentionalism, 6, 82–88
- Intentions
 - speaker intentions, 50, 55–79, 85, 91, 109
- Interpretive interaction, 82, 88–94

L

- Logic
 - and context dependence, 4
 - first-order, 121, 125–146, 187
 - of indexicals and demonstratives, 16
- Logical truths, 98, 128–130, 133–135, 137, 146, 195
- Logophoricity, 154, 163

M

- Meaning
 - linguistic meaning, 2, 3, 10, 81, 115, 196, 207–214, 218, 220, 227–232, 234, 237–240, 242, 247, 278
- Meaning rules, 9, 205, 215–220
- Metasemantics
 - coordination account metasemantics, 55–58, 60, 68, 74, 75, 79
 - direct, 6, 31, 33, 39–41, 58
 - indirect, 5, 6, 30–37, 39–43, 47–50, 58, 67–69, 72, 75
 - objective, 57–60, 62, 63, 65–67, 74, 75, 79

Minimalism, 11, 273, 276, 291, 292

Modals

- deontic modals, 198
- epistemic modals, 298, 299, 301, 304, 307, 308, 311, 312

Monsters, 8, 9, 114, 150, 152, 153, 157, 163–176, 179–203

O

- Occasion-sensitivity, 227–242
- Operators, 8, 9, 114, 118, 127, 149–153, 155, 160, 161, 165, 166, 174, 175, 177, 179–181, 184, 185, 189, 190, 193–195, 201, 297, 304, 308, 310, 311

P

- Perlocutionary intentions, 10, 252–258, 260
- Pragmatic resolution, 213, 300–304, 314
- Pragmatics
 - and semantics, 3, 10, 182, 212, 227
 - truth conditional pragmatics, 9, 10, 205–223, 227–229, 231–233, 291
- Presuppositions, 3, 4, 6, 33–35, 37, 42, 47, 56, 59, 60, 62–64, 66, 67, 71, 72, 74, 82, 87, 172–174, 193, 274
- Pronouns, 30, 43, 56, 58, 81, 125, 126, 146, 150, 153, 154, 159, 161–163, 165–167, 171, 207, 255, 308
- Proof theory, 131, 144
- Proper names, 8, 24, 151, 159, 161, 182, 183, 207
- Proposition
 - singular proposition, 24, 26
 - structured proposition, 4, 127, 152, 167, 181, 184, 228, 234
 - See also* Content
- Propositional attitude, 153, 176, 189, 190, 266

R

- Reference frame, 7, 100–107, 109
- Reference time (RT), 7, 100–111
- Reflexive-referential, 5, 16–25, 27
- Relativism
 - moderate relativism, 11, 284, 288, 289
- Relevance, 9, 93, 154, 237, 246, 251, 281

S

Saliency, 85, 219, 220, 222, 284
 Scare-quoting, 7, 114, 119–121, 123
 Self-ascription, 26, 27, 157, 159
 Semantics
 event semantics, 308
 formal, 9, 136, 187–194, 197, 205, 206,
 208, 211–213, 215
 and pragmatics, 182, 212
 truth conditional, 206, 210, 227, 229, 231,
 292
 Semantic underdeterminacy, 235
 Semiotics, 180, 182, 185, 186, 201
 Situated unenriched illocutions, 285
 Speech act
 subsential speech acts, 11, 283–292
 Supplementives, 6, 40, 55–59, 66–68, 71, 72

T

Temporal indexicals, 7, 102–104, 155
 Tenses, 7, 15, 43, 56, 99, 127, 171, 190, 267,
 308

Travis cases, 229–232, 234–240
 Truth, 5, 19, 34, 60, 94, 98, 115, 126, 150, 180,
 210, 229, 265, 287, 310

Truth conditions

 referential truth conditions, 5, 19–22, 24
 reflexive truth conditions, 5, 19–22, 24,
 286, 287

U

Unarticulated constituents, 4, 11, 268, 285,
 288–291
 Utterance
 distributed utterance, 7, 8, 113–123
 Utterance time (UT), 7, 100–111, 310,
 311

V

Vagueness, 34, 210