Government Surveillance and Internet Search Behavior

April 29, 2015

Abstract

This paper uses data from eleven countries on the search volume of select keywords from before and after the surveillance revelations of June 2013, to analyze whether Google users' search behavior changed as a result. The surveillance revelations are treated as an exogenous shock in information about how closely users' internet searches were being monitored by the US government. Each search term was independently rated for its degree of privacy sensitivity along multiple dimensions. Using panel data, our results suggest that cross-nationally, users were less likely to search using search terms that they believed might get them in trouble with the US government. In the US, this was the main subset of search terms that were affected. However, internationally there was also a drop in traffic for search terms that were rated as personally sensitive. Our results suggest that there is a chilling effect on search behavior from government surveillance on the Internet, and that government surveillance programs may damage the international competitiveness of US-based internet firms.

1 Introduction

On June 6, 2013, new information began to emerge about the surveillance practices of the US government, starting with the publication of leaked classified documents by then-Guardian columnist Glenn Greenwald. These contained revelations about the 'PRISM' program, which is a codename for what appears to be a mass electronic surveillance data mining program managed by the National Security Agency (NSA). The NSA's slides disclosed partnerships of a kind with nine major tech companies, including Microsoft, Google, Yahoo!, AOL, Skype and others, for the NSA to obtain real-time data content.

The revelations provoked a highly public and ongoing controversy, both from domestic privacy activists and from international governments concerned about the privacy of their own citizens. What is not clear is how actual user online behavior changed as a result of the controversy. Broad surveys of US residents report some ambivalence about the program. An initial Pew survey conducted in July 2013 suggested that 50% of US citizens approved of the government phone metadata and Internet data surveillance programs disclosed to that point, and 44% disapproved of them;[1] in a later Pew survey from January 2014, the proportion disapproving had risen to 53%. A November 2013 survey by the US writers' organization PEN shows 28% of its responding members as having self-censored in response to the surveillance revelations.[2] On the firm side, Castro (2013) discusses a survey conducted by the Cloud Security Alliance which showed 56 percent of non-US members said that they would be less likely to use a US-based cloud computing service as a consequence of the PRISM revelations.

Unlike this survey-based data already in the public domain, our study aims to be the first reasonably comprehensive empirical study to document whether and how actual user behavior, in terms of the use of search engines, changed after the surveillance revelations

[1] http://www.people-press.org/2013/07/26/few-see-adequate-limits-on-nsa-surveillance-program/
[2] http://www.pen.org/sites/default/files/Chilling\%20Effects_PEN\%20American.pdf

began. We examine whether search traffic for more privacy-sensitive search terms fell after the exogenous shock of publicity surrounding the NSA's activities. To be clear, we are not measuring responses to the phenomenon of mass government surveillance *per se*. Such surveillance has been conducted for a long time, with varying levels of public scrutiny and concern. We instead measure the effects of such surveillance activities becoming much more widely known and understood.

To explore this question, we collected data on internet search term volume before and after June 6, 2013, to see whether the number of searches was affected by the PRISM revelations. We collected this data using Google Trends, a publicly available data source which has been used in other studies to predict economic and health behaviors (Choi and Varian, 2012; Carneiro and Mylonakis, 2009). We collected data on the volume of searches for the US and its top ten international trading partners (in order, Canada, China, Mexico, Japan, Germany, South Korea, the United Kingdom, France, Brazil and Saudi Arabia) during all of 2013 for 282 search terms.

These 282 search terms came from three different sources: a Department of Homeland Security list of search terms it tracks on social media sites (DHS (2011), pp. 20-23), Google's top 50 search terms for 2013 (Google, 2013), and a crowd-sourcing exercise to identify potentially embarrassing search terms that did not implicate homeland security.

These sources are clearly non-random and are intended to provide an external source of search terms to study. Having obtained this list, we then employed independent raters to rank these search terms in terms of how likely their usage was to get the user in trouble with the US government or with a 'friend.' We make this distinction between trouble with the government and trouble with a friend in the ratings to try and tease apart the potential for differences in behavioral responses to privacy concerns emanating from the personal domain and the public domain. There are different policy implications if people self-censor searches that may be used to identify potentially criminal behavior and those which are just personally

sensitive. We use these ratings as moderators in our empirical analysis to understand the differential effects of the revelations on different search terms.

We find that the Google Trends search index fell for "high government trouble" search terms by roughly a 10% after the Snowden revelations. The fact we observe any significant effect in the data is surprising, given skepticism about whether the surveillance revelations were capable of affecting search traffic at such a macro level in the countries concerned. There was a smaller but still significant decline for search terms that raters thought gave them an above average likelihood of getting in trouble with a friend. This was driven, however, by US international trading partners rather than by the US itself. We check the robustness of these results in a variety of ways, including using different time windows as a falsification check and using controls for news coverage. We also use a variety of other proxy measures for privacy concerns to check the robustness of our results.

This paper aims to contribute to two strands of the academic literature.

The first is an economic literature that aims to measure demand for privacy. Acquisti et al. (2013) and Brandimarte et al. (2012) use behavioral economics to study what affects consumer preferences for privacy. Gross and Acquisti (2005) examine demand for privacy settings on a social network. Goldfarb and Tucker (2012) use refusals to volunteer private information as a proxy measure for privacy demand, to study inter-generational shifts in privacy demand. Since we differentiate between user behavior in eleven different countries, we are able to compare quantitatively the reactions of users in those different countries to the same exogenous shock revealing the collection of their search data by the US government, and therefore to assess in a novel manner the demand in those countries for privacy in their search terms.

The second literature measures the effect on consumer behavior of government privacy policies and practices and their implications for commercial outcomes. Miller and Tucker (2009); Adjerid et al. (2015) have shown mixed effects of privacy regulations on the diffusion

of digital health. Romanosky et al. (2008) show mixed effects for data breach notification laws on identify theft, while Goldfarb and Tucker (2011); Campbell et al. (2015) document potentially negative effects of privacy regulation for the competitiveness of digital advertising. To our knowledge, there is little empirical research which uses observed behavior to investigate how the policies of governments towards surveillance affect consumer behavior and commercial outcomes.[3]

2 Data and Background

2.1 PRISM Revelations

On June 6, 2013, new information emerged about the surveillance practices of the US government, starting with the publication of leaked classified documents by Guardian columnist Glenn Greenwald.[4] These contained revelations about the 'PRISM' program, which is a codename for what appears to be a mass electronic surveillance data mining program managed by the National Security Agency (NSA). The NSA's slides disclosed partnerships of a kind with nine major tech companies, including Microsoft, Google, Yahoo!, AOL, Skype and others, for the NSA to obtain real-time data content.

[3]There have however, been anecdotal accounts. For example, `http://www.usatoday.com/story/tech/2014/02/27/nsa-resistant-products-obama-tech-companies-encryption-overseas/5290553/` quotes Andrew Jaquith, chief technology officer at cloud-security firm SilverSky. as saying 'Suspicion of U.S. vendors is running at an all-time high.'

[4]On the morning of June 6, 2013, the 'Verizon scandal' also disclosed to the public that phone companies including Verizon had been ordered by a secret court to continuously disclose the metadata associated with all calls - location, caller, callee and call duration - subject to a routine renewal every 90 days. Though we believe that the PRISM revelations are likely to have a more direct causal mechanism when it comes to search engine behavior, we acknowledge that the multiplicity of revelations on the same date means that we cannot separately identify the effect of the PRISM and Verizon revelations. We also acknowledge that since this date, many further scandals have resulted from the same set of leaked documents. However, it seems appropriate to study the impact of the revelations as a whole, and therefore to begin at the point of initial disclosure on June 6. Later information also suggested that the NSA might itself, on its disclosed slides, have been overstating the official nature of its partnerships with the companies named. Further disclosures at later dates relating to other programs, including XKEYSCORE and TEMPORA, could also, for highly informed users, have further affected their search behavior. However, as our study considers the impact on search behavior among the general public of the publicization of surveillance, rather than the unpublicized operation of the programs themselves, we believe these fine-grained distinctions are not material for our analysis.

The US government emphasized in its initial response that the 'authority [under which the program falls] was created by the Congress and has been widely known and publicly discussed.' (DNI, 2013), but it was not generally understood prior to June 2013 that the authority in question, Section 702 of the FISA Amendments Act of 2008, authorized consumer data held by such companies, including data on US individuals' search behavior, to be made available to the US government on a mass rather than an individualized basis.[5]

2.2 Data

The data we use is derived from Google Trends, which is a public source of cross-national search volume for particular search terms. Prior to collecting this data, we had to identify a list of search terms which would provide appropriate and reasonable coverage of the kind of search terms that may have been affected by the PRISM revelation, and also a quasi-control set of search terms. We use search terms from three sources: A DHS list, a crowdsourced "embarrassing terms" list, and a third list of the terms Google Trends itself defines as the "top search terms" for 2013.

We use search terms from a 2011 US government list (DHS, 2011) of "suspicious" selectors that might lead to a particular user being flagged for analysis by the NSA. This is a 2011 list provided for the use of analysts working in the Media Monitoring Capability section of the National Operations Center, an agency under the Department of Homeland Security. The list was made public in 2012, and continued to be used and reproduced within DHS up to the time of the surveillance revelations (DHS, 2013); as far as we are aware, it remains in effect. It is therefore the most relevant publicly available document for assessing the kinds of search terms which the US government might be interested in collecting under PRISM or under its other programs aimed at gathering Google search data, even though it is focused

[5]Freedom of Information Act litigation brought by privacy organization EPIC in 2013-14 would, had it been successful, have required the release of the Office of Legal Counsel memos containing the interpretation of Section 702 that authorizes collection under PRISM, but an adverse ruling means that these memos are still secret. See EPIC v. DOJ, 2013 DC No. 1:13-cv-01848 (BAH), accessed at `https://epic.org/foia/doj/olc/prism` on April 14, 2015.

on surveillance of social media websites rather than search engines. The full list is in the appendix as Tables 12 and 13.

Our overall aim in establishing a reasonable list of separate personally 'embarrassing' search terms was to find terms that would not implicate national security issues of interest to DHS, or duplicate any term found in that list, but which would still plausibly cause personal embarrassment if third parties found that you had been searching on them.[6] We crowdsourced this list for this purpose using a group of participants in the Cambridge Co-Working Center, a startup incubator located in Cambridge, MA. The participants were young (20s-30s), well-educated, and balanced equally between men and women. The full list of 101 search terms presented in Tables 14 and Table 15 in the appendix is the result of that crowd-sourcing process.

We also wanted to obtain a list of more "neutral" search terms to use as a quasi-control. To find this, we turned to Google itself, which releases an annual list of what it describes as the 'Zeitgeist' - the top search terms people searched for during that year. This is clearly not a raw list, and is not accompanied with data on absolute search volume. It is limited in two ways. First, the list dates to mid-December of 2013, so omits the second half of December. Second, the list is curated: It contains no offensive or obscene search terms. However, it is a reasonable snapshot of the variety of search terms of most interest to the world's users of the Google search engine. Google makes only the top ten search terms easily available, but it does make public, in a harder-to-access form, all of the terms occupying places #11 to #100 on their list. We identified all 100, but used only the top 50 search terms from the list. This was because the list does not exclude politically-charged searches, and therefore included 'Edward Snowden' at #97, which would defeat the purpose of using this list for external controls. We reproduce the full list in the appendix as Table 16; Google's Zeitgeist interface is accessible at `www.google.com/zeitgeist`.

[6] We instructed the group to not include obscenities or words relating to obscene acts.

We then collected data on the weekly search volume for each of our 282 search terms from Google Trends.[7] We collected data separately on the volume of searches for the US and its top ten international trading partners (in order, Canada, China, Mexico, Japan, Germany, South Korea, the United Kingdom, France, Brazil and Saudi Arabia) according to the US Census.[8] This led to a dataset of 161,876 observations.

Google Trends data has been used in a variety of academic studies to measure how many people are searching for specific items in order to better inform economic and even health forecasting (Choi and Varian, 2012; Carneiro and Mylonakis, 2009). The methodology behind Google Trends is somewhat opaque. Google states that 'Google Trends analyzes a percentage of Google web searches to determine how many searches have been done for the terms you have entered compared to the total number of Google searches done during that time.' Google also says it excludes duplicate searches and searches made by a few people. The key disadvantage of the Google Trends data from our perspective is that Google only provides the data in a normalized format. Google states, 'Normalized means that sets of search data are divided by a common variable, like total searches, to cancel out the variable's effect on the data.' This means that a typical Google index for a search term in a region spans 0-100. Theoretically, this does not affect the validity of the directional nature of our results. The key issues come from the fact that the data is not provided in terms of absolute number of searches, making it harder to project economic outcomes or enumerate the actual changes to searches. However, as there are no alternative data providers of clickstream data that provide sufficient international scope, we decided to accept this limitation.[9]

[7] `www.google.com/trends`

[8] `http://www.census.gov/foreign-trade/statistics/highlights/top/top1312yr.html`

[9] One possible response of users to the surveillance revelations would have been to switch from using a PRISM-implicated search engine such as Google's to an encrypted search engine, of which the best-known examples are DuckDuckGo and Tor. Though DuckDuckGo usage certainly increased, it was from such a low base that by the end of 2013 DuckDuckGo traffic worldwide represented 0.4% of Google's traffic (see http://www.theguardian.com/technology/2014/jan/09/anonymous-search-tool-duckduckgo-1bn-queries-2013-google). Similarly, the user base of Tor quintupled during the summer and fall of 2013, from under 1 million daily users to somewhat over 5 million daily users, before settling back down to a steady base

Table 1 provides summary statistics of the distribution of the different search terms and weekly search volume in our Google Trends data. These summary statistics apply to the 2013 data we focus on in our analysis, but we also collected other years of data that we use in subsequent falsification checks.

Table 1: Summary Statistics for Google Trends Data

	Mean	Std Dev	Min	Max	Observations
Search Volume	13.3	21.0	0	100	161876
Crowd-Sourced Embarrassing Term	0.34	0.47	0	1	161876
DHS Sensitive Search Term	0.48	0.50	0	1	161876
Google Top 50 Search Term	0.18	0.38	0	1	161876
United States	0.091	0.29	0	1	161876
After Prism Revelations	0.58	0.49	0	1	161876
Number of News Stories	51.9	188.6	0	2313	161876

In our analysis, we focus on data on searches on Google, simply due to cross-national data availability. Google remains the world's dominant search engine, with a January 2014 worldwide market share of over 70%.[10] Table 2 uses data from the NSA's PRISM slides on the dates major search engines began to participate in the PRISM program.[11] The three major US search firms - Microsoft, Yahoo! and Google - are listed as the first three participants, and by the time of the surveillance revelations of 2013 had been involved with the program

of around 2-3 million daily users worldwide during early 2014. At its peak, it represented approximately one five-hundredth of Google's daily search volume (see https://metrics.torproject.org/ for Tor usage statistics for the relevant timeframe). The overall proportion of Web traffic that is encrypted has also risen sharply, presumably in response to the surveillance revelations (see http://www.wired.com/2014/05/sandvine-report/), but that rise does not affect Google Trends' data.

[10]There are specific countries in our dataset where Google's presence in the national market substantially differs from this average. For example, in China Google was not one of the top two search providers during 2013, and those wishing to use Google Search in China often do so using a VPN which would lead to their search results being attributed to a different country. In South Korea, Google also has a minor share of the market, and in Japan it takes second place with a market share of 40%. In the other seven countries in our analysis, Google enjoyed during 2013 a dominant market position. See for further details http://returnonnow.com/internet-marketing-resources/2013-search-engine-market-share-by-country/.

[11]The extent to which their participation has been active or passive, and the extent to which senior decision makers at these firms were aware of the firms' "participation" in PRISM, is still unclear, and is expected to be clarified in the course of ongoing litigation.

for approximately six, five and four years respectively.

Table 2: PRISM Data Collection Providers

Provider Name	PRISM Data Collection Start Date
Microsoft	September 2007
Yahoo!	March 2008
Google	January 2009
Facebook	June 2009
PalTalk	Dec 2009
YouTube	December 2010
Skype	February 2011
AOL	March 2011
Apple	October 2012

Source: `http://www.washingtonpost.com/wp-srv/special/politics/prism-collection-documents/`

Though we tried to collect a set of search terms from a set of diverse sources in order to span the idea of searches that may be viewed as neutral, personally sensitive or government sensitive, it is not clear how an average user would view the privacy sensitivity of each search term. For example, the DHS list of search terms contains phrases such as "agriculture" which may not be commonly viewed as a search term which would get you into trouble with the government or something that the government may be tracking.[12] Furthermore, some phrases could be both personally sensitive and sensitive in the eyes of the government. For example, a search term like 'marijuana legalization' may be personally embarrassing if friends did not know you used the drug, and may also be viewed as a search phrase that could lead to trouble with the government.

To address this shortcoming and the variation within each list to which each search term presented a privacy threat, we collected further data to try and establish externally which of these search terms reflected politically and personally sensitive topics. We asked close to 6,000 workers on Amazon Mechanical Turk to evaluate a single search term each.

[12]We may reasonably infer that the US government was monitoring this particular term out of concern about terrorist attacks on the agricultural supply chain, but the phrase by itself is not evocative of terrorist threats.

Similar crowdsourcing techniques have been used by Ghose et al. (2012) to design rankings for search results. Recent research into the composition of workers on Mechanical Turk has suggested that in general they are reliable and representative for use as subjects in psychological experiments (Paolacci et al., 2010; Buhrmester et al., 2011). However, we recognize that in demographics they are likely to skew younger than the average population (Tucker, 2015).

In the survey, we asked participants to rate a term by how likely it is that it would 'get them into trouble' or 'embarrass' them with their family, their close friends, or with the US government. We also asked them to rate how privacy-sensitive they considered the term, how much they would like to keep the search secret, and how likely they would be to try and delete their search history after using this term. Table 17 in the appendix reproduces the survey questions in full. All ratings used a five-point Likert scale, where 1 reflects the least 'sensitive' and 5 reflects the most 'sensitive' rating. Table 3 reports the results of this extra step in our search term evaluation process. As might be expected, the terms on the DHS list are most likely to be rated as 'getting you in trouble with the US government', at a mean value of 1.62 out of 5; though overall the DHS terms are not on average rated close to the highest value possible of 5 on the scale because they contain many apparently innocuous terms, such as "symptoms" and "agriculture." The search terms from the 'embarrassing' list were rated the most likely to embarrass the user to their family or close friends, at mean values of between 2.2 and 2.3 out of 5 in terms of whether they would embarrass the user if their close friends or family knew about them, whether the user would want to keep the search secret or delete their search history, but at a lower sensitivity value of 1.59 in terms of whether the search would get them into trouble with the U. S. government. The Google Trend terms were, as expected, generally rated the least embarrassing, with mean sensitivity values ranging between 1.24 and 1.43 out of 5 on all measures. Table 18 in the appendix presents the cross-index correlations.

Table 3: 'Privacy' Rating of Google Search Terms by Source

	DHS Term Mean	Embarrassing Term Mean	Google Top Search Mean	Total Mean
Trouble Employer	1.57	1.87	1.38	1.64
Trouble Family	1.42	1.71	1.24	1.49
Trouble Friend	1.41	1.64	1.25	1.46
Trouble Government	1.62	1.59	1.25	1.55
Embarrassed Employer	1.66	2.42	1.54	1.90
Embarrassed Family	1.53	2.27	1.42	1.76
Embarrassed Friend	1.49	2.21	1.43	1.73
Embarrassed Government	1.63	1.82	1.28	1.63
Keep Search Secret	1.65	2.28	1.42	1.83
Privacy-Sensitive Rating	1.68	2.30	1.38	1.84
Delete Search History	1.66	2.28	1.41	1.83

3 Empirical Analysis

3.1 Model-Free Analysis

Before turning to econometric analysis, we present some 'model-free' evidence about major trends in the data in graph form.

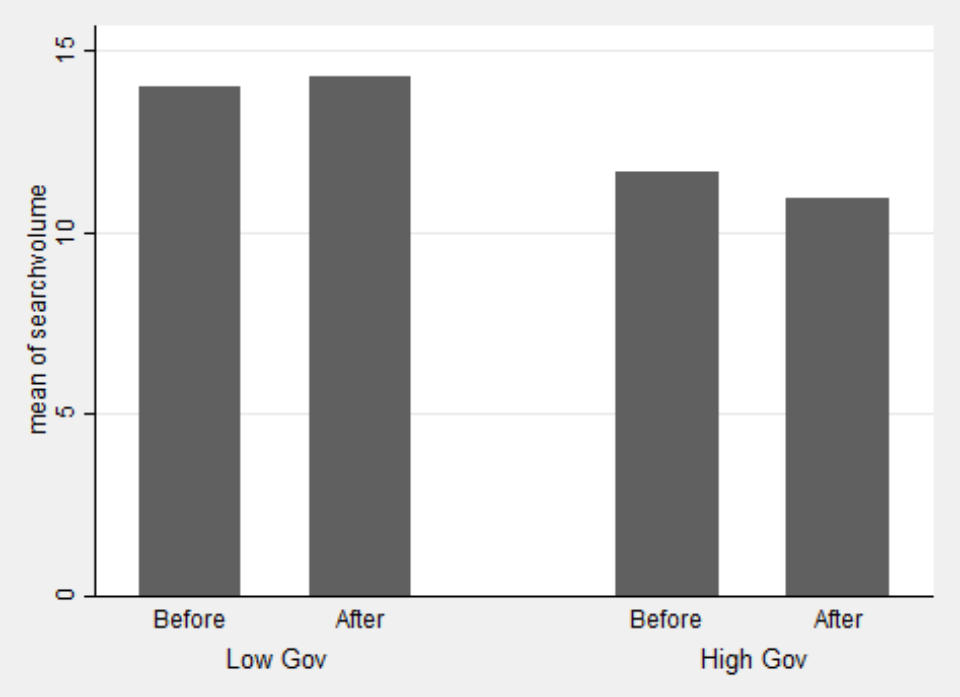

Figure 1: Search Volume before and after PRISM Revelations

Figure 1 presents our initial analysis where we separate out aggregate search volume for 2013 before and after the revelations and by whether that search term was rated as above average in terms of causing trouble for the searcher with the US government. Overall, across the eleven countries we study, search terms that were rated as being unlikely to get you in trouble with the US government exhibited a slight rise in traffic. However, search terms that were rated as being more likely to get you in trouble with the US government exhibited a distinct fall in traffic.

It is important to be cautious in interpreting the changes in these bar charts. At such

an aggregate level it is hard to straightforwardly assert that the changes we observe were attributable to the surveillance (and particularly the PRISM) revelations.

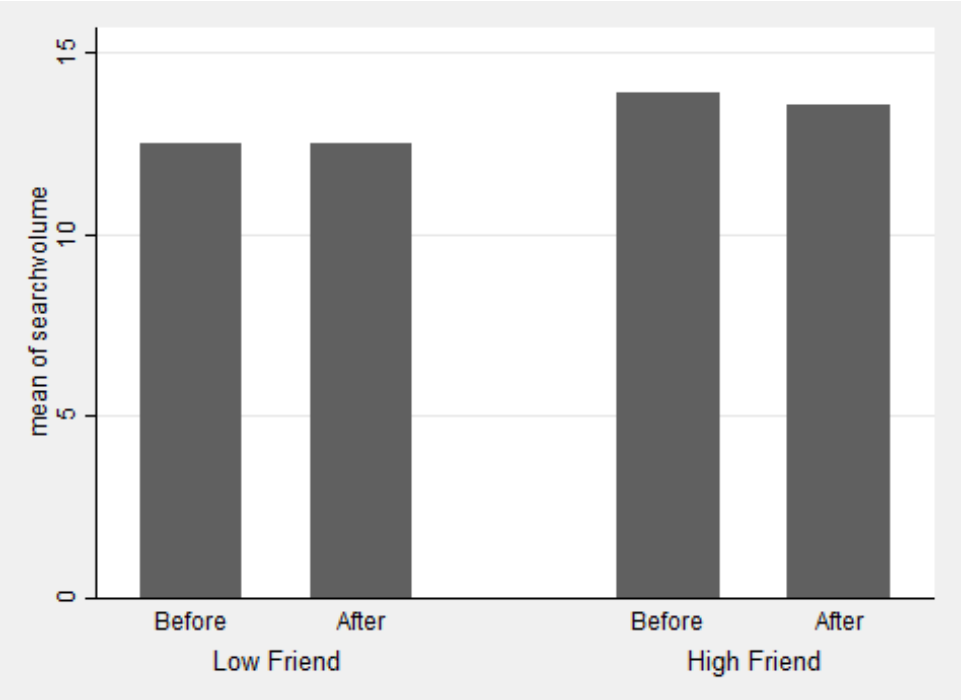

Figure 2: Search Volume before and after PRISM Revelations

Next, we reran this analysis to compare search traffic in the eleven countries using terms that were rated as having a low level of likelihood that it would lead the user to be in trouble if a close friend knew about the user's search ("low-friend"), versus terms that had a high level ("high-friend"). As shown by Figure 2, the overall pattern more or less holds: traffic for low-friend terms holds steady, and traffic for high-friend terms falls, though by an amount that is less pronounced than in Figure 1.

3.2 Econometric Analysis

The empirical analysis is straightforward. We compare before and after the PRISM revelations with multiple different controls in a panel data setting to see whether there were measurable shifts in the patterns of search behavior after the revelations relative to before. This kind of approach has been described as 'regression discontinuity' in Busse et al. (2006), which examines changes around a short time window surrounding a policy change. However, we recognize that in papers which use the exact timing of a particular event as their discontinuity, rather than some arbitrary exogenous threshold, identification is always going to be weaker than in a more standard regression discontinuity (Hahn et al., 2001).

We model the search volume rate $SearchVolume_{ijt}$ for search term i in country j on week t in the following manner:

$$\begin{aligned} SearchVolume_{ijt} = \beta PrivacySensitivity_i \times PostPrism_t \\ +\gamma_i + \theta_j + \delta_t + \epsilon_i \end{aligned} \tag{3.1}$$

γ is a series of fixed effects for each of our 282 keywords, θ_j is a series of fixed effects for each country, and δ_t is a series of fixed effects at the weekly level. The fixed effects γ control for the different natural levels of search volume for each of the different search terms. θ_j captures general differences in search volume across countries. δ_t captures week-by-week variation in search term volume that may be driven by work patterns or holidays. This means that our major coefficient of interest is β, which measures the differential in search volume for keywords that were more sensitive for this measure after the PRISM revelations. The main effect of $PostPrism$ is collinear with the weekly fixed effects and is consequently dropped from our regressions. Similarly, the main effect of our index of sensitivity, $PrivacySensitivity_i$, is collinear with the keyword fixed effects and is consequently

dropped from the regression.

Table 4 presents our initial results. The first three columns focus on a specification where we use a binary indicator to mark whether a search term was considered to be above average in terms of its likelihood to lead to trouble with the government or trouble with a friend. Column (1) presents the results for all countries.[13] The results suggest that search terms rated as having an above-average likelihood of getting the searcher in trouble with the government, fell two index points relative to an average index of 13 points. Search terms rated as having an above-average likelihood of getting the searcher in trouble with a friend fell by half an index point. Column (2) presents results for the US only; Column (3) presents results for the ten non-US countries in our study. We see that US-based search traffic falls by quite a large extent in the Google index for terms that are perceived as having an above-average likelihood of getting you in trouble with the US government, whereas non-US traffic also falls but by a smaller magnitude. However, by contrast, in non-US countries there is a significant fall in the volume of search terms that are perceived as having an above-average likelihood of getting the searcher in trouble with a friend.

The second three columns of Table 4 present a complementary specification where, rather than using an indicator variable, we use the full scale for how likely raters perceived this search term as leading to trouble with a government or friend. The results are similar to those in the non-parametric specification in columns (1)-(3) and in subsequent regressions we use these specifications, as they use more of our available data.

[13]Results for a specification which allowed each country to have an individual weekly time-trend produced almost identical results, presumably because of the way that Google creates its index.

Table 4: Initial Results

	All Countries	US	Non-US	All Countries	US	Non-US
	(1)	(2)	(3)	(4)	(5)	(6)
Post Prism × High Gov Trouble	-2.068***	-3.280***	-1.947***			
	(0.171)	(0.351)	(0.170)			
Post Prism × High Friend Trouble	-0.432**	0.353	-0.510***			
	(0.179)	(0.357)	(0.177)			
Post Prism × Gov Trouble				-1.219***	-2.227***	-1.118***
				(0.195)	(0.343)	(0.196)
Post Prism × Friend Trouble				-1.459***	-0.114	-1.594***
				(0.275)	(0.439)	(0.277)
Country Fixed Effects	Yes	No	Yes	Yes	No	Yes
Keyword Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Week Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	161876	14716	147160	161876	14716	147160
R-Squared	0.510	0.859	0.509	0.510	0.859	0.509

OLS Estimates. Dependent Variable Is Search Volume Index As Reported By Google Trends over 2013.
Robust Standard Errors Clustered At Search Term Level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
The main effects of $PostPrism$ and the two $Sensitive_i$ terms are collinear with the week and keyword fixed effects and consequently both terms are dropped from the regression.

Overall, these results provide empirical evidence that the surveillance revelations caused a substantial chilling effect relating to users' willingness to enter search terms that raters considered would get you into trouble with the US government. We also see that outside the US there is a large and significant drop in the search terms which are more likely to get you into trouble if a friend found out you used that search term. This suggests that international users, in contrast to US users, reduced their relative number of searches for personally sensitive terms.s

3.2.1 Robustness

A natural concern is whether other factors could plausibly have shifted user behavior in early June relating to these specific keywords. However, the keywords cover a large variety of topics, so another news story relating to a small portion of them, such as an extreme weather event (for the DHS search terms) or the release of a new movie (for the Google Zeitgeist terms) is unlikely to have shifted behavior for the whole. A more plausible concern would be Google-specific or search-engine specific, i.e. whether there was an internal change in the way the search engine operated at the time that would have a coherent and similar effect on search behavior in multiple countries.

To address this and tie the effect more closely to the actual PRISM revelations, we tried to establish whether our finding was robust to a narrower time window, so we reran the analysis using only data from five weeks before and five weeks after the first surveillance revelations on June 6, 2013. This is reported as Table 5. As we might expect, there is a similar fall in search traffic for high-trouble terms in the US in the shorter period, but what is notable is that the effect for the non-US results is smaller for the shorter period, though still evenly distributed across search terms that are sensitive both in the governmental domain and the personal domain. This would fit with an interpretation that the short-term shock in the country where the scandals originated was larger, but that over time the worldwide

Table 5: Robustness to 10 Week Interval

	All Countries	US	Non-US
	(1)	(2)	(3)
Post Prism × Gov Trouble	-0.980**	-2.000**	-0.878*
	(0.459)	(0.974)	(0.461)
Post Prism × Friend Trouble	-1.085*	-0.851	-1.109*
	(0.603)	(1.281)	(0.606)
Country Fixed Effects	Yes	No	Yes
Keyword Fixed Effects	Yes	Yes	Yes
Week Fixed Effects	Yes	Yes	Yes
Observations	34243	3113	31130
R-Squared	0.523	0.875	0.523

OLS Estimates. Dependent Variable Is Search Volume Index As Reported By Google Trends for 5 weeks before and 5 weeks after PRISM Revelation date.
Robust Standard Errors Clustered At Search Term-Level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
The main effects of $PostPrism$ and the two $Sensitive_i$ terms are collinear with the week and keyword fixed effects and consequently both terms are dropped from the regression.

ramifications of the surveillance revelations became more apparent, and user search behavior shifted accordingly. In general, this check reassures us that unrelated shocks to search behavior and search engines, like the UK government's attempt to block certain types of pornography at the end of 2013, are not driving our results.

Table 6: Falsification Test for 2012

	All Countries	US	Non-US
	(1)	(2)	(3)
Post Prism × Gov Trouble	-0.159	0.0702	-0.182
	(0.164)	(0.237)	(0.162)
Post Prism × Friend Trouble	0.237	0.202	0.241
	(0.244)	(0.388)	(0.243)
Country Fixed Effects	Yes	No	Yes
Keyword Fixed Effects	Yes	Yes	Yes
Week Fixed Effects	Yes	Yes	Yes
Observations	164989	14999	149990
R-Squared	0.585	0.916	0.587

OLS Estimates. Dependent Variable Is Search Volume Index As Reported By Google Trends for 2012. The placebo Post-PRISM date in this table is June 4, 2012.
Robust Standard Errors Clustered At Search Term Level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
The main effects of $PostPrism$ and $Sensitive_i$ are collinear with the week and keyword fixed effects and consequently both terms are dropped from the regression.

We also tried to rule out seasonality as being a driver of our results by repeating the analysis of Table 4 for 2012, using exactly the same June date. Table 6 reports the results of this falsification check. All the coefficients are reassuringly insignificant. This suggests that it is not seasonality brought about by comparing quarters 1 and 2 with quarters 3 and 4 that is driving our results.

Table 7 explores another concern, which is that rather than be associated directly with the content revelations, the effects we measure are simply a function of news coverage. To explore this, we gathered data from Factiva on the number of news stories in each country which mentioned the NSA and Edward Snowden. We use this data as a proxy for how extensive news coverage was in that country and in that week. Table 7 shows our results which reflect this additional robustness check. Our earlier results hold, suggesting that the change we measure is not media-driven. In general, news coverage seems to be negatively related to overall search volume.

Table 7: Controlling for the Effect of News Coverage

	All Countries	US	Non-US
	(1)	(2)	(3)
Post Prism × Gov Trouble	-1.219***	-2.227***	-1.118***
	(0.195)	(0.343)	(0.196)
Post Prism × Friend Trouble	-1.459***	-0.114	-1.594***
	(0.275)	(0.439)	(0.277)
Number of News Stories	-0.00103***		-0.00140**
	(0.000313)		(0.000697)
Country Fixed Effects	Yes	No	Yes
Keyword Fixed Effects	Yes	Yes	Yes
Week Fixed Effects	Yes	Yes	Yes
Observations	161876	14716	147160
R-Squared	0.510	0.859	0.509

OLS Estimates. Dependent Variable Is Search Volume Index As Reported By Google Trends over 2013. Robust Standard Errors Clustered At Search Term Level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The main effects of $PostPrism$ and the two $Sensitive_i$ terms are collinear with the week and keyword fixed effects and consequently both terms are dropped from the regression. The news story variable for the US-only result is collinear with the week fixed effects and is also dropped from the regression.

Another concern is that our findings might be an artifact of the particular sensitivity factors we decided to focus on; that is, whether the person felt that the use of such a search

term might get them into trouble with either the government or a friend. We chose this distinction as it was a clear contrast between the personal and governmental domain when it came to privacy sensitivity, but we wanted to check that there was not something about those particular questions which drove our results, for example, whether the use of the word 'trouble' was particularly emotive.

Table 8 shows the robustness of our findings to alternative ways of measuring for privacy sensitivity for all countries. The summary statistics for these alternative measures are presented in Table 3 and the full text of the questions is noted in Table 17. Columns (1) and (2) reflect the likelihood of trouble with a family member or employer. While the family member question echoes strongly the 'trouble with a friend' question that we focus on in this paper, the 'trouble with employer' represents an intermediate step between the personal and more public domain of privacy. As discussed by Acquisti and Fong (2013), an employer's relationship with a employee and use of personal data to shape that relationship is a new challenge for privacy policy in the internet era. Columns (3)-(6) use an alternative measure which is a scale of how likely the search was to cause 'embarrassment' in turn with a friend, family member, employer or government. In general, the results persist using this alternative language to the word 'trouble' which is our primary focus in our regressions. However, the effect is less precisely measured, reflecting that embarrassment is perhaps a lesser fear among our raters than actual trouble, which meant that there was less strong distinctions between the ratings of the different words. Column (7) uses a more straightforward and direct measure of actual 'privacy sensitivity' in terms of our raters. Again the results hold, though the point estimate is smaller, perhaps reflecting ambiguity in how privacy as a concept was interpreted among our raters. The final two columns (8) and (9) use more concrete behavioral rather than perceptual measures, asking raters in turn for the likelihood they would actively keep a search with this search term secret, or delete the search. Our results hold with these behavioral rather than perceptual measures.

Table 8: Robustness to Other Privacy Measures (Pooled Regression for All Countries)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Trouble Family× Post Prism	-1.597***								
	(0.176)								
Trouble Employer× Post Prism		-1.407***							
		(0.143)							
Embarrassed Employer× Post Prism			-0.916***						
			(0.114)						
Embarrassed Family× Post Prism				-0.765***					
				(0.116)					
Embarrassed Friend× Post Prism					-0.752***				
					(0.122)				
Embarrassed Government× Post Prism						-2.078***			
						(0.145)			
Privacy-Sensitive Rating× Post Prism							-1.113***		
							(0.114)		
Keep Search Secret× Post Prism								-0.965***	
								(0.116)	
Delete Search History× Post Prism									-1.081***
									(0.116)
Country Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Keyword Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Week Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	161876	161876	161876	161876	161876	161876	161876	161876	161876
R-Squared	0.509	0.509	0.509	0.509	0.509	0.510	0.509	0.509	0.509

OLS Estimates. Dependent Variable Is Search Volume Index As Reported By Google Trends over 2013.
Robust Standard Errors Clustered At Search Term Level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
The main effects of $PostPrism$ and the various $Sensitive_i$ terms are collinear with the week and keyword fixed effects and consequently both terms are dropped from the regression.

Table 9: Robustness to Other Privacy Measures (U.S. Only)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Trouble Family× Post Prism	-1.249***								
	(0.291)								
Trouble Employer× Post Prism		-1.217***							
		(0.243)							
Embarrassed Employer× Post Prism			-0.682***						
			(0.183)						
Embarrassed Family× Post Prism				-0.301*					
				(0.177)					
Embarrassed Friend× Post Prism					-0.324*				
					(0.186)				
Embarrassed Government× Post Prism						-2.171***			
						(0.271)			
Privacy-Sensitive Rating× Post Prism							-0.985***		
							(0.184)		
Keep Search Secret× Post Prism								-0.790***	
								(0.186)	
Delete Search History× Post Prism									-0.890***
									(0.190)
Keyword Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Week Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	14716	14716	14716	14716	14716	14716	14716	14716	14716
R-Squared	0.859	0.859	0.859	0.859	0.859	0.859	0.859	0.859	0.859

OLS Estimates. Dependent Variable Is Search Volume Index As Reported By Google Trends over 2013.
Robust Standard Errors Clustered At Search Term Level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
The main effects of $PostPrism$ and the various $Sensitive_i$ terms are collinear with the week and keyword fixed effects and consequently both terms are dropped from the regression.

Table 10: Robustness to Other Privacy Measures (Non U.S. Countries)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Trouble Family× Post Prism	-1.632***								
	(0.178)								
Trouble Employer× Post Prism		-1.426***							
		(0.147)							
Embarrassed Employer× Post Prism			-0.940***						
			(0.116)						
Embarrassed Family× Post Prism				-0.812***					
				(0.117)					
Embarrassed Friend× Post Prism					-0.795***				
					(0.122)				
Embarrassed Government× Post Prism						-2.068***			
						(0.147)			
Privacy-Sensitive Rating× Post Prism							-1.126***		
							(0.116)		
Keep Search Secret× Post Prism								-0.983***	
								(0.118)	
Delete Search History× Post Prism									-1.100***
									(0.118)
Country Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Keyword Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Week Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	147160	147160	147160	147160	147160	147160	147160	147160	147160
R-Squared	0.508	0.508	0.508	0.508	0.508	0.509	0.509	0.508	0.508

OLS Estimates. Dependent Variable Is Search Volume Index As Reported By Google Trends over 2013.
Robust Standard Errors Clustered At Search Term Level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
The main effects of $PostPrism$ and the various $Sensitive_i$ terms are collinear with the week and keyword fixed effects and consequently both terms are dropped from the regression.

Table 9 shows the same robustness for the US alone. Table 10 shows the same robustness check for all non-US countries. A comparison of Table 9 and Table 10 echoes some of the patterns seen in Table 4 where other countries prove more sensitive to private domains of privacy sensitivity than to the governmental domain.

4 Mechanism

To explore a potential underlying mechanism that could explain the disparity between the US and elsewhere, we divided the countries in our dataset into different categories depending on their familiarity with and attitudes towards surveillance. We measured this in three ways. First, we relied on a Pew study which was part of the 'Pew Global Attitudes Project' which asked a sample of that country's citizens 'Is American monitoring of Your Country's citizens acceptable or unacceptable?' We then divided up the countries into whether they had above-average or below-average acceptance of US monitoring practices. This is intended to proxy for the extent to which the citizens of that country are accepting of surveillance. Second, we identified two ways of categorizing countries by how familiar its citizens were with surveillance. We divided up our sample and compared the results for countries which have a history of monitoring their citizens' online searches and those that do not. This categorization was done on the basis of the work of the Citizen Lab Internet research group who uses computer servers to scan for the distinctive signature of technology which enables surveillance by governments.[14] As a secondary measure we also divided up our sample by whether or not that country had a system of ID cards providing the government with a universal identifier for citizens.

Table 11 reports the results of these splits. In general, the point estimates of the coefficient on 'government trouble' are, in absolute terms, larger for countries where there is more familiarity and acceptance of surveillance. The point estimates of the coefficient on 'friend

[14] http://www.nytimes.com/2013/01/16/business/rights-group-reports-on-abuses-of-surveillance-and-censo
html?_r=0

trouble' are, in absolute terms, smaller relatively for these countries.

We caution, however, that as we were are relying on stratification among 10 countries, this evidence is more suggestive than statistically robust; we do not have sufficient statistical power to distinguish sufficiently between the countries. However, with this caveat, these results might suggest that what we are measuring overall in terms of the drop in search volume related to search terms which may get the searcher in trouble with the US government, is a pragmatic response of a citizenry accustomed to or accepting of government surveillance. However, the point estimate on the coefficient on 'friend trouble' is larger, in absolute terms, in countries which have lower experience and acceptance of surveillance. This might suggest that the response that we measure internationally on more personally sensitive search terms is less pragmatic but instead a more instinctive response to the realization that search behavior is not private.

Table 11: Familiarity with Surveillance

	Prism Unacceptable	Prism Acceptence	Low Own Surveillance	High Own Surveillance	No ID Card	ID Card
	(1)	(2)	(3)	(4)	(5)	(6)
Post Prism × Gov Trouble	-1.181***	-1.238***	-1.110***	-1.150***	-0.997***	-1.239***
	(0.295)	(0.278)	(0.215)	(0.362)	(0.251)	(0.271)
Post Prism × Friend Trouble	-1.667***	-1.564***	-1.625***	-1.467***	-1.633***	-1.554***
	(0.385)	(0.412)	(0.300)	(0.547)	(0.362)	(0.375)
Country Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Keyword Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Week Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	58864	58864	117728	29432	73580	73580
R-Squared	0.557	0.558	0.534	0.645	0.579	0.507

OLS Estimates. Dependent Variable Is Search Volume Index As Reported By Google Trends over 2013.
Robust Standard Errors Clustered At Search Term Level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

5 Conclusion

This study is the first to provide substantial empirical documentation of a chilling effect, both domestically in the shorter term and internationally in the longer term, that appears to be related to increased awareness of government surveillance online. Furthermore, this chilling effect appears in countries other than the US to apply to search behavior that is not strictly related to the government but instead forms part of the private domain.

Our findings have the following policy implications. From an economic perspective, our finding that there was an effect on international Google users' browsing behavior has potential policy implications for the effects of government surveillance on international commerce. From a US competitive standpoint, the longer-run effect observed on international Google users' search behavior indicates that knowledge of US government surveillance of Google could indeed affect their behavior. At the most limited end of the spectrum, it could steer them away from conducting certain searches on US search engines; at the most severe end of the spectrum, they might choose to use non-US search engines. Such effects may not be limited simply to search engines. For example, as Google's services are embedded in a large array of products, it could potentially hinder sales of Android-enabled mobile phones. Though preliminary attempts are being made to work towards initial measures of the economic impact of surveillance revelations (Dinev et al., 2008), no systematic study yet exists. All we can do, within the context of our data, is to indicate that on the basis of the effects we find, the strong possibility of substantial economic effects exists, and to suggest that such potential adverse economic impacts should be incorporated into the thinking of policy makers regarding the appropriateness of mass surveillance programs.

There are limitations to the generalizability of our findings. First, we are not sure how the results generalize outside of the search domain towards important tech industries such as the rapidly growing US cloud computing industry. Second, we are not sure how the revelations

affected search on Google's major competitors, such as Bing and Yahoo! Search. It may be that the effect on their services was lessened by reduced media focus on them relative to Google in the light of the PRISM revelations and potentially the extent to which users anticipated that their servers may be located outside of the US. Third, our results are focused on the effects of revelations about government surveillance as opposed to the direct effects of government surveillance *per se*. Notwithstanding these limitations, we believe that our study provides an important first step in understanding the potential for effects of government surveillance practices on commercial outcomes and international competitiveness.

References

Acquisti, A. and C. M. Fong (2013). An experiment in hiring discrimination via online social networks. *Available at SSRN 2031979*.

Acquisti, A., L. K. John, and G. Loewenstein (2013). What is privacy worth? *The Journal of Legal Studies 42*(2), 249 – 274.

Adjerid, I., A. Acquisti, R. Padman, R. Telang, and J. Adler-Milstein (2015). The impact of privacy regulation on technology adoption: The case of health information exchanges. *Management Science*.

Brandimarte, L., A. Acquisti, and G. Loewenstein (2012). Misplaced confidences: Privacy and the control paradox. *Social Psychological and Personality Science*.

Buhrmester, M., T. Kwang, and S. D. Gosling (2011). Amazon's Mechanical Turk, a new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science 6*(1), 3–5.

Busse, M., J. Silva-Risso, and F. Zettelmeyer (2006). $1,000 cash back: The pass-through of auto manufacturer promotions. *American Economic Review 96*(4), 1253–1270.

Campbell, J. D., A. Goldfarb, and C. Tucker (2015). Privacy Regulation and Market Structure. *Forthcoming, Journal of Economics & Management Strategy*.

Carneiro, H. A. and E. Mylonakis (2009). Google trends: a web-based tool for real-time surveillance of disease outbreaks. *Clinical infectious diseases 49*(10), 1557–1564.

Castro, D. (2013, August). How much will PRISM cost the U.S. cloud computing industry?

Choi, H. and H. Varian (2012). Predicting the present with Google trends. *Economic Record 88*(s1), 2–9.

DHS (2011). National Operations Center Media Monitoring Capability Analyst's Desktop Binder. *Department of Homeland Security*. Accessed from `http://www.scribd.com/doc/82701103/Analyst-Desktop-Binder-REDACTED` on February 28, 2014.

DHS (2013). Privacy Impact Assessment for the Office of Operations Coordination and Planning: Publicly available social media monitoring and situational awareness initiative update, Appendix B: Terms used by the NOC MMC when monitoring social media sites. *Department of Homeland Security Website*.

Dinev, T., P. Hart, and M. R. Mullen (2008). Internet privacy concerns and beliefs about government surveillance an empirical investigation. *The Journal of Strategic Information Systems 17*(3), 214 – 233.

DNI (2013, June 8). Facts on the collection of intelligence pursuant to section 702 of the Foreign Intelligence Surveillance Act. *Director of National Intelligence*. Accessed at `http://www.dni.gov/files/documents/Facts\%20on\%20the\%20Collection\%20of\%20Intelligence\%20Pursuant\%20to\%20Section\%20702.pdfonFebruary28,2014.`

Ghose, A., P. G. Ipeirotis, and B. Li (May/June 2012). Designing ranking systems for hotels on travel search engines by mining user-generated and crowdsourced content. *Marketing Science 31*(3), 493–520.

Goldfarb, A. and C. Tucker (2012). Shifts in privacy concerns. *American Economic Review: Papers and Proceedings 102*(3), 349–53.

Goldfarb, A. and C. E. Tucker (2011). Privacy regulation and online advertising. *Management Science 57*(1), 57–71.

Google (2013). Google Zeitgeist 2013 report. *Google*.

Gross, R. and A. Acquisti (2005). Information revelation and privacy in online social networks. In *Proceedings of the 2005 ACM workshop on Privacy in the electronic society*, WPES '05, New York, NY, USA, pp. 71–80. ACM.

Hahn, J., P. Todd, and W. Van der Klaauw (2001). Identification and estimation of treatment effects with a regression-discontinuity design. *Econometrica 69*(1), 201–209.

Miller, A. and C. Tucker (2009, July). Privacy protection and technology adoption: The case of electronic medical records. *Management Science 55*(7), 1077–1093.

Paolacci, G., J. Chandler, and P. G. Ipeirotis (2010). Running experiments on Amazon Mechanical turk. *Judgment and Decision making 5*(5), 411–419.

Romanosky, S., R. Telang, and A. Acquisti (2008). Do Data Breach Disclosure Laws Reduce Identity Theft? *Mimeo, Carnegie Mellon*.

Tucker, C. E. (2015). The reach and persuasiveness of viral video ads. *Marketing Science 34*(2), 281–296.

A Search Terms in Full

Table 12: DHS Search Terms

	Gov Trouble Rating
DHS	1.55
TSA	1.35
UCIS	1.50
agent	1.10
agriculture	1.05
air marshal	1.74
alcohol tobacco and firearms	2
anthrax	2.76
antiviral	1.65
assassination	2.44
authorities	1.35
avian	1.24
bacteria	1.15
biological	1.25
border patrol	1.37
breach	1.63
burn	1.63
center for disease control	1.60
central intelligence agency	1.55
chemical	2.10
chemical agent	2.21
chemical burn	1.85
chemical spill	1.89
cloud	1.05
coast guard	1.30
contamination	1.70
cops	1.39
crash	1.22
customs and border protection	1.65
deaths	1.25
dirty bomb	3.74
disaster assistance	1.37
disaster management	1
disaster medical assistance te	1.18
dndo	1.84
domestic security	2.15
drill	1.06
drug administration	1.79
drug enforcement agency	1.85
ebola	1.17
emergency landing	1.42
emergency management	1.76
emergency response	1.40
epidemic	1.68
evacuation	1.35
explosion	2.20
explosion explosive	3.15
exposure	1.50
federal aviation administratio	1.10
federal bureau of investigatio	1.63
first responder	1
flu	1.58
food poisoning	1.60
foot and mouth	1.45
fusion center	1.75
gangs	1.56
gas	1.55
h1n1	1.44
h5n1	1.60
hazardous	1.61
hazmat	1.35
homeland defense	1.42
homeland security	1.75
hostage	2.06
human to animal	2.20
human to human	1.45
immigration customs enforcemen	1.47
incident	1.47
infection	1.60
Total	1.62

Table 13: DHS Search Terms

	Gov Trouble Rating
influenza	1.20
infrastructure security	1.75
law enforcement	1.30
leak	1.40
listeria	1.47
lockdown	1.70
looting	2.11
militia	1.89
mitigation	1.45
mutation	1.58
national guard	1.37
national laboratory	1.45
national preparedness	1.60
national security	1.79
nerve agent	3.21
north korea	1.75
nuclear	2.10
nuclear facility	2.42
nuclear threat	2.17
organized crime	2.32
outbreak	1.60
pandemic	1.42
pipe bomb	4
plague	1.68
plume	1.11
police	1.20
pork	1.16
powder white	2.30
prevention	1.15
public health	1.30
quarantine	2.15
radiation	1.85
radioactive	2.05
recall	1.39
recovery	1.30
red cross	1.20
resistant	1.50
response	1.10
ricin	2.60
riot	1.60
salmonella	1.26
sarin	2.89
screening	1.30
secret service	1.89
secure border initiative	1.55
security	1.21
shooting	1.90
shots fired	2.11
sick	1.10
small pox	1.79
spillover	1.11
standoff	1.47
state of emergency	1.40
strain	1.39
swat	1.55
swine	1.25
symptoms	1
tamiflu	1.50
task force	1.15
threat	1.70
toxic	1.44
tuberculosis	1.20
united nations	1.20
vaccine	1.20
virus	1.40
wave	1.05
world health organization	1.22
Total	1.63

Table 14: Embarrassing Search Terms

	Friend Trouble Rating
abortion	2.30
accutane	1.26
acne	1.10
adultery	2.26
agenda 21	1.47
aids	1.63
alcoholics anonymous	2.11
alien abduction	1.40
animal rights	1.16
anonymous	1.18
atheism	1.45
bail bonds	1.55
bankruptcy	2
bittorrent	1.37
black panthers	1.60
body odor	1.63
breathalyzer	1.65
casinos	1.21
celebrity news	1.11
chemtrails	1.78
coming out	2.05
communism	1.37
conspiracy	1.37
cop block	1.35
cutting	2.75
debt consolidation	1.79
depression	2
divorce lawyer	1.65
drones	1.42
eating disorder	2
erectile dysfunction	2
escorts	2.60
feminism	1.11
filesharing	1.45
fireworks	1.20
food not bombs	1.45
gay rights	1.47
gender reassignment	2.11
ghosts	1.25
gulf of tonkin	1.32
guns	2.05
herpes	1.89
hitler	1.85
hoarding	1.45
honey boo boo	1.33
incontinence	1.45
islam	1.25
keystone	1.16
kkk	2.11
Total	1.62

Table 15: Embarrassing Search Terms

	Friend Trouble Rating
larp	1.74
liposuction	1.26
lolcats	1.16
lonely	1.68
lost cause	1.26
marijuana legalization	1.50
marx	1.42
my little pony	1.50
nickelback	1.85
nose job	1.60
occupy	1.70
online dating	2
pest control	1.17
peta	1.20
police brutality	1.25
polyamory	1.80
porn	1.95
pregnant	1.70
protest	1.61
psychics	1.65
revolution	1.40
sexual addiction	2.45
shrink	1.65
socialism	1.22
sovereign citizen	1.21
sperm donation	2.06
strip club	2.26
suicide	2.68
tampons	1.85
tax avoidance	1.90
therapist	1.45
thrush	1.17
torrent	1.28
transhumanism	1.47
turner diaries	1.74
tuskegee	1.16
unions	1.28
vaccines and autism	1.33
vegan	1.30
viagra	2.16
warts	1.55
weed	2.11
weight loss	1.50
white power	3.05
white pride	2.47
wicca	1.80
witchcraft	1.84
world of warcraft	1.35
Total	1.66

Table 16: Google Search Terms

	Friend Trouble Rating
2014 fifa world cup	1.30
aaron hernandez	1.53
adrian peterson	1.11
amanda bynes	1.45
atari breakout	1.16
boston marathon	1.32
charlie hunnam	1.35
chennai express	1.20
cory monteith	1
despicable me 2	1.20
django unchained	1.37
government shutdown	1.35
gravity	1.25
gta 5 cheats	1.65
harlem shake	1.37
hugo chavez	1.32
ios 7	1
iphone 5s	1.10
iron man 3	1.17
james gandolfini	1.32
jennifer lawrence	1.15
jodi arias	1.42
kevin ware	1.42
kim kardashian baby	1
lea michele	1.26
les miserables	1.30
lou reed	1.05
malala yousafzai	1.58
man of steel	1.11
miley cyrus vmas	1.16
mindy mccready	1.11
nelson mandela	1.15
nexus 5	1.26
north korea	1.35
oscar pistorius	1.50
pacific rim	1.20
paul walker	1.21
paula deen	1.22
playstation 4	1.11
red sox	1.26
robin thicke	1.20
royal baby	1.16
salve jorge	1.68
samsung galaxy s4	1.25
the conjuring	1.05
trayvon martin	1.37
typhoon	1
windows 81	1.11
world war z	1.21
xbox one	1.05
Total	1.25

Table 17: Survey Questions in Full

How likely is it that you would be in trouble if the US government found out you used this search term?
How likely is it that you would be in trouble if your employer found out you used this search term?
How likely is it that you would be in trouble if a family member found out you used this search term?
How likely is it that you would be in trouble if a close friend found out you used this search term?
How likely is it that you would feel embarrassed if the US government found out you used this search term?
How likely is it that you would feel embarrassed if your employer found out you used this search term?
How likely is it that you would feel embarrassed if a family member found out you used this search term?
How likely is it that you would feel embarrassed if a close friend found out you used this search term?
How likely are you to view this search term as privacy-sensitive?
How likely is it that you would like to keep this search term secret?
How likely is it that you would delete the search history on your computer after using this search term?

B Correlation of Privacy Indices

Table 18: Cross-correlation table

Variables	delete_search	employer_embarrassed_searc	employer_trouble_search	family_embarrassed_search	family_trouble_search	friend_embarrassed_search	friend_trouble_search
delete_search	1.00						
employer_embarrassed_searc	0.69	1.00					
employer_trouble_search	0.56	0.75	1.00				
family_embarrassed_search	0.74	0.83	0.62	1.00			
family_trouble_search	0.61	0.63	0.73	0.73	1.00		
friend_embarrassed_search	0.72	0.80	0.58	0.90	0.67	1.00	
friend_trouble_search	0.62	0.57	0.66	0.66	0.87	0.65	1.00
government_embarrassed_sea	0.63	0.69	0.67	0.64	0.63	0.62	0.62
government_trouble_search	0.50	0.56	0.72	0.50	0.65	0.48	0.64
privacy_search	0.86	0.70	0.57	0.74	0.61	0.71	0.64
secret_search	0.87	0.70	0.56	0.76	0.61	0.74	0.63