# Government Surveillance and Internet Search Behavior

Alex Marthews*
and Catherine Tucker‡

February 17, 2017

**Abstract**

This paper displays data from the US and its top 40 trading partners 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 search terms that were deemed both personally-sensitive and government-sensitive were most negatively affected by the PRISM revelations, highlighting the interplay between privacy concerns relating to both the government and the private individual. Perhaps surprisingly, the largest 'chilling effects' were not found in countries conventionally treated as intelligence targets by the US, but instead in countries that were more likely to be considered allies of the US. We show that this was driven in part by a fall in searches on health-related terms. Suppressing health information searches potentially harms the health of search engine users and, by reducing traffic on easy-to-monetize queries, also harms search engines' bottom line. In general, 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 profitability of US-based internet firms relative to non-US-based internet firms.

**Keywords: surveillance, Snowden, privacy, PRISM, chilling effects, search engines, international trade**

**JEL Classification: D12, D78, E65, F14, H56, M38**

*Digital Fourth, Cambridge, MA.

†MIT Sloan School of Management, MIT, Cambridge MA and National Bureau of Economic Research.

‡We thank participants at the 2014 Future of Privacy & Data Security Regulation Roundtable, the 2014 Privacy Law Scholars Conference and the 2016 Hackers on Planet Earth conference for useful comments. We acknowledge funding from the NSF, Law and Economics program at George Mason University and the Vanguard Charitable Foundation. All errors are our own.

# 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 in the British ‘Guardian’ newspaper. These contained revelations about the ‘PRISM’ program, 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 about US citizens.

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 reporting 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 saying 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

[1] Pew Research Center, “Few See Adequate Limits on NSA Surveillance Program,But More Approve than Disapprove”, July 26, 2013, available at http://www.people-press.org/2013/07/26/few-see-adequate-limits-on-nsa-surveillance-program/, accessed February 17, 2017.

[2] “Chilling Effects: NSA Surveillance Drives US Writers to Self-Censor”, PEN American Center, November 12, 2013; available at https://www.pen.org/sites/default/files/2014-08-01_Full\%20Report_Chilling\%20Effects\%20w\%20Color\%20cover-UPDATED.pdf, accessed February 17, 2017.

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 began.[3] 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.

In general, after news spread of what the documents showed, there was much press discussion about whether the revelations would in fact affect user behavior. On the one hand, the revelations were of a nature that it might be intuitive to expect some change in user search behavior within the US, and perhaps also in countries already known to be major targets of US foreign surveillance, relating to search terms that they expected would be likely to get them in trouble with the US government, such as, say, 'pipe bomb' or 'anthrax.' On the other hand, the argument was also made that people were, or ought already to have been, aware that the US government conducted surveillance on the Internet, and that they might therefore already have 'baked in' an expectation of such surveillance into their behavior, making a new effect as a result of these revelations unlikely to be observed (Cohen, 2013). Last, it is not clear that even if people express concerns that their privacy has been intruded upon, actual behavioral change will result. It is therefore an empirical question to determine whether there were in fact such behavioral changes.

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

[3]Though subsequent research papers (Penney, 2016; Cooper, 2017) have reused aspects of our methodology, it is still reasonable to characterize our study as the first to apply empirical techniques to the study of the actual impact of surveillance on citizen behavior.

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 40 international trading partners during all of 2013 for 245 search terms.

These 245 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); a neutral list of search terms based on the most common local businesses in the US; and a crowd-sourcing exercise to identify potentially embarrassing search terms that did not implicate homeland security.

These sources are obviously 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 to 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 users self-censor searches that they believe may signal potentially criminal behavior, versus if users self-censor searches that are personally sensitive without any criminal implications. We use these ratings as moderators in our empirical analysis to understand the different effects of the revelations on different search terms.

We find that the Google Trends search index fell, for search terms that were deemed troubling from both a personal and private perspective, by roughly 4% after the revelations. 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 then show that internationally, the effect was stronger in countries where English is the first language. We also show that the effect was stronger in countries where surveillance was less acceptable and citizens were less used to surveillance by their government. Perhaps surprisingly, we found that the largest 'chilling' effects were not found in countries traditionally considered intelligence targets by the US, but instead in countries that were more likely to be considered allies of the US.

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. First, there is an entire literature on political ignorance and apathy (Somin, 2016), suggesting that broadly speaking, individuals are poorly informed about political matters and have few incentives to become better informed. This scandal could be expected to generate behavioral changes among a minority of politically engaged people, but, given the low level of information on the part of the public about surveillance matters, it might easily be considered unlikely to generate meaningful behavioral change beyond that limited audience. Second, the lack of empirical proof of chilling effects has been a topic of significant discussion in legal academia,[4] so for this audience the very idea of a study that is able to measure such effects is neither straightforward or intuitive.

This paper aims to contribute to three 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 41 different countries, we are able to compare quantitatively the reactions of users in those different countries to the

[4]See, for example (Richards, 2013), published immediately before the Snowden revelations, which argues that though the chilling effects of surveillance are 'empirically unsupported, [...] such criticisms miss the point. The doctrines encapsulated by the chilling effect reflect the substantive value judgment that First Amendment values are too important to require scrupulous proof to vindicate them.'

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) and 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 identity 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 using observed behavior to investigate how the policies of governments towards surveillance affect consumer behavior and commercial outcomes.

The third literature we contribute to is on the privacy paradox. Those who have found a privacy paradox (Gross and Acquisti, 2005; Barnes, 2006; Athey et al., 2017) identify that people in practice, when faced with short-term decisions, do not change their information sharing habits or are not willing to pay even a small amount for the preservation of the privacy that they articulate as an important value to them; and that similarly, if a service is offered to them that is privacy-compromising but free, most will opt for it over a service that carries a fee but that does not compromise privacy. Here, we see that in the actual usage of a free service, people will shape their searches in order to avoid surveillance.

# 2 Data

## 2.1 Search Engine Data

Table 1 uses data from the NSA's PRISM slides on the dates major search engines began to participate in the PRISM program.[5] 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 for approximately six, five and four years respectively.

Table 1: 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/`

The data we use is derived from Google Trends, which is a public source of cross-national search volume for particular search terms. We focus on data on searches on Google, simply due to international data availability. Google remains the world's largest search engine, with a market share of around 70% at the time of the PRISM revelations. We exploit variation in the size of its presence in subsequent regressions cross-nationally where we explore differences in consumer behavior in countries where Google's search engine presence is less sizable.

Google Trends data has been used in a variety of academic studies to measure how many

[5] 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.

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. To do this, each data point is divided by the total searches of the geography and time range it represents, to compare relative popularity. The resulting numbers are then scaled to a range of 0 to 100.'[6] 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 search volumes. However, as there are no alternative data providers of clickstream data that provide sufficient international scope, we decided to accept this limitation.

## 2.2 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 baseline searches for common local businesses and services.

We use search terms from a US government list (DHS, 2011) of "suspicious" selectors

[6] `https://support.google.com/trends/answer/4365533?hl=en`

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 on surveillance of social media websites rather than search engines. The full list is in the appendix as Tables A-1 and A-2.

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.[7] 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 A-3 and A-4 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.We emphasize that our use of the term 'quasi-control' does not mean that our specification should be thought of as a classic difference-in-difference. Instead, this more neutral set of search terms should be thought of as simply a group of searches that were plausibly treated less intensively by the revelations about PRISM.

To find a more neutral set of search terms we turned to the nature of Google as a search

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

engine.[8] Users across the world use Google to search for local services and businesses. This type of search behavior provides a reasonable baseline measure of usage of search engines. To obtain words to capture this behavior, we first obtained a list of the most common local businesses in the US based on the North American Industry Classification System.[9] We associated this list with search terms that would plausibly capture these businesses.[10]

We then collected data on the weekly search volume for each of our 245 search terms from Google Trends.[11] We collected data separately on the volume of searches for the US and its top 40 international trading partners according to the IMF.[12] The top ten in order are Canada, China, Mexico, Japan, Germany, South Korea, the United Kingdom, France, Brazil and Saudi Arabia. The remaining 30 are Argentina, Australia, Austria, Belgium, Colombia, Denmark, Egypt, Hong Kong (treated separately from China), India, Indonesia, Iran, Israel, Italy, Malaysia, the Netherlands, Nigeria, Norway, Pakistan, the Philippines, Poland, Russia, Singapore, South Africa, Spain, Sweden, Switzerland, Taiwan, Thailand, Turkey and the United Arab Emirates. This led to a dataset of 523,340 observations on the week-country-search term level.

Table 2 provides summary statistics of the distribution of the different search terms and weekly search volume in our Google Trends data. The value of 0.396 for 'Crowd-Sourced

[8] In earlier versions of this paper, we used data from Google Zeitgeist (`www.google.com/zeitgeist`) as a source of potentially neutral words. Since that earlier version, we have greatly expanded the list of countries we study, rendering Zeitgeist no longer a satisfactory set of controls, because so much of it focused on US cultural figures such as American football player Aaron Hernandez. This tended to provide a very uneven baseline of search behavior internationally.

[9] Fitness and Recreational Sports Centers (NAICS: 71394), Full-Service Restaurants (72211), Homes for the Elderly (62331), All Other Amusement and Recreation Industries (71399), Used Merchandise Stores (45331), Meat Processed from Carcasses (31161), Landscape Architectural Services (54132), Beauty Salons (81211), Carpet and Upholstery Cleaning Services (56174), and Child Day Care Service (62441).

[10] Most categories were straightforward and captured by the search terms: Gym, restaurant, nursing home, thrift store, butcher, gardener, beauty salon, cleaners, and childcare. For the Amusement and Recreation industry, we included arcade, movies and weather to capture searches an individual might perform related to recreation.

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

[12] IMF World Economic Outlook Database, available at `https://www.imf.org/external/pubs/ft/weo/2016/02/weodata/index.aspx`, accessed February 16, 2017.

Embarrassing Term' indicates that the crowd-sourced embarrassing terms comprise 39.6% of the dataset. Similarly, the value .555 for 'DHS Sensitive Search Term' indicates that DHS terms comprise 55.5% of the dataset. These summary statistics apply to the 2013 data we focus on in our analysis, but we also collected data from 2012 that we use in subsequent falsification checks.

Table 2: Summary Statistics for Google Trends Data

| | Mean | Std Dev | Min | Max | Observations |
|---|---|---|---|---|---|
| Search Volume | 10.19 | 15.4 | 0 | 100 | 522340 |
| Crowd-Sourced Embarrassing Term | 0.396 | 0.49 | 0 | 1 | 522340 |
| DHS Sensitive Search Term | 0.555 | 0.50 | 0 | 1 | 522340 |
| Neutral | 0.0490 | 0.22 | 0 | 1 | 522340 |
| United States | 0.0244 | 0.15 | 0 | 1 | 522340 |
| After Prism Revelations | 0.577 | 0.49 | 0 | 1 | 522340 |
| Number of News Stories | 18.57 | 105.5 | 0 | 2313 | 522340 |

## 2.3 Sensitivity of Search Terms

Though we tried to collect search terms from a diverse set of sources, in order to obtain a reasonable range of search terms that were 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 as something that the government may be tracking.[13] 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 took support for legalization as evidence that you used the drug, and may also be viewed as a

[13] 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.

search term that could lead to trouble with the US government given marijuana's continued illegal status under federal law.

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 5,000 workers on Amazon Mechanical Turk to evaluate a single search term each. Each of our 246 keywords was rated by 20 different Mechanical Turkers.

We set a qualification level such that each worker had to have a 'Hit Approval Rate (%),' which is the proportion of tasks they have performed in the task that were approved by the employer, of greater than 95%, to try to further assure the quality of the workers we recruited. As it turned out, none of our workers had an approval rating of less than 100%.

We also checked to see if our ratings were altered if we removed workers who took a shorter or longer time than usual, but did not see any significant effects.

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' with their family, their close friends, or with the US government.[14] Table 3 reproduces the survey questions we study in this paper. All ratings used a five-point Likert scale, where 1 reflects the least 'sensitive' and 5 reflects the most 'sensitive' rating. Table 4 reports the results of this extra step in our search term evaluation process. As might be

[14]We also asked them to rate how privacy-sensitive or embarrassing 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. In earlier versions of the paper we showed robustness to using these alternative metrics.

Table 3: Survey Questions Wording

| |
|---|
| 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? |

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 "cloud" and "incident." The search terms from the 'embarrassing' list were rated at a lower sensitivity value of 1.59 in terms of whether the search would get them into trouble with the U. S. government, but at 1.64 in terms of getting you in trouble with a friend. The local business terms, which are intended to be neutral, were, as expected, generally rated the least embarrassing, with mean sensitivity values ranging between 1.04 and 1.11 out of 5 on all measures. Table A-6 in the appendix presents cross-index correlations.

Table 4: 'Trouble' Rating of Google Search Terms by Source

| | DHS Term Mean | Embarrassing Term Mean | Neutral Mean | Total Mean |
|---|---|---|---|---|
| Trouble Employer | 1.57 | 1.87 | 1.11 | 1.67 |
| Trouble Family | 1.42 | 1.71 | 1.06 | 1.52 |
| Trouble Friend | 1.41 | 1.64 | 1.04 | 1.49 |
| Trouble Government | 1.62 | 1.59 | 1.04 | 1.58 |

## 2.4 Pre-trends in Data

In our data analysis we treat the PRISM revelations as having occurred on June 6, 2013.[15] 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.[16] Therefore we treat the PRISM revelations as an exogenous shock to how informed search engines users were about the extent to which the US government was monitoring their search behavior.

One concern, of course, is whether before the PRISM revelations the search volume for the different terms were moving in a similar direction. To explore this, we constructed a figure that explored the extent to which search terms of different types moved in parallel prior to the revelations.

Figure 1 shows the pre-trends for each of the categories of keywords we study. They show

[15]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 Upstream, 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.

[16]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.

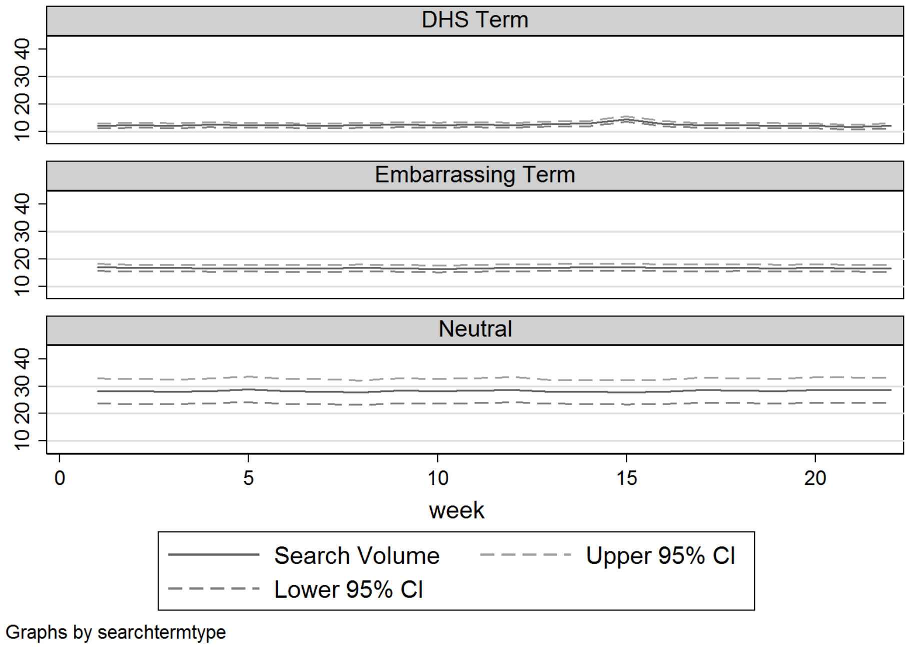


Figure 1: Evidence of Common Trends Prior to the PRISM Revelations

similar trends.[17]

One worry is of course that the sensitivity of these metrics changed over the period we study. To evaluate this, we repeated the ratings exercise two years after the initial Mechanical Turk measurement exercise in the US, and observed an extremely high correlation between the two measurement exercises - with a Spearman correlation of 0.95 - and a raw correlation of 0.88. We also tried running our regression excluding the few search terms whose sensitivity had changed during the time period - for example, celebrities such as 'Honey Boo Boo' who are no longer as famous as they were in 2013. Our results remained the same.

[17]The only notable exception is an uptick in searches for DHS terms in April 2013. This appears to have been the result of the Boston Marathon bombing, as people searched for information about bombs. As a result of this uptick, we ran a robustness study where we excluded April 2013 from our data, and obtained similar results.

# 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.

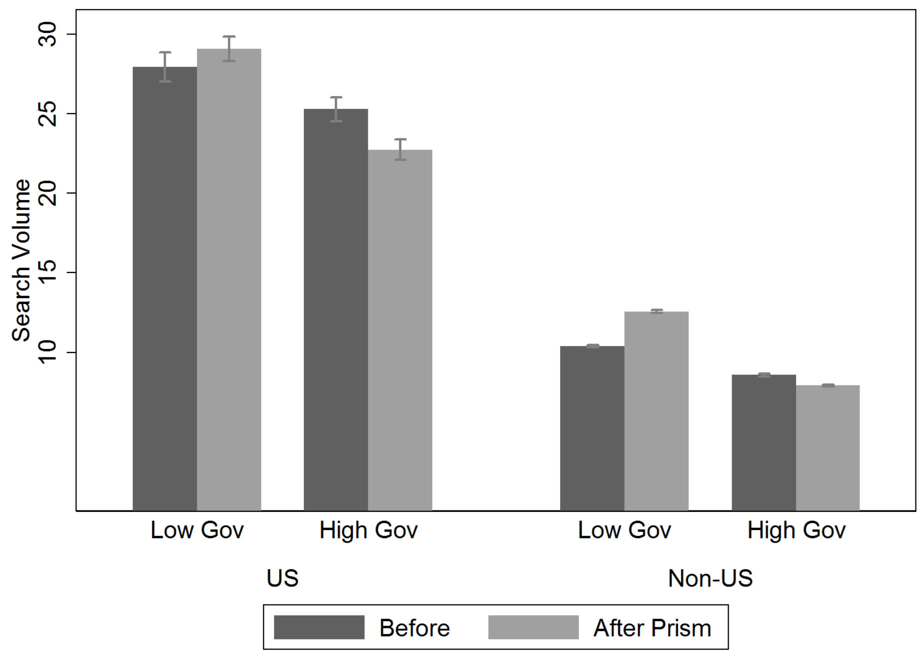


Figure 2: Search Volume Before and After PRISM Revelations

Figure 2 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-median in terms of causing trouble for the searcher with the US government. Overall, across the 41 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, particularly in the US.

Next, we reran this analysis to compare search traffic in the countries using terms that

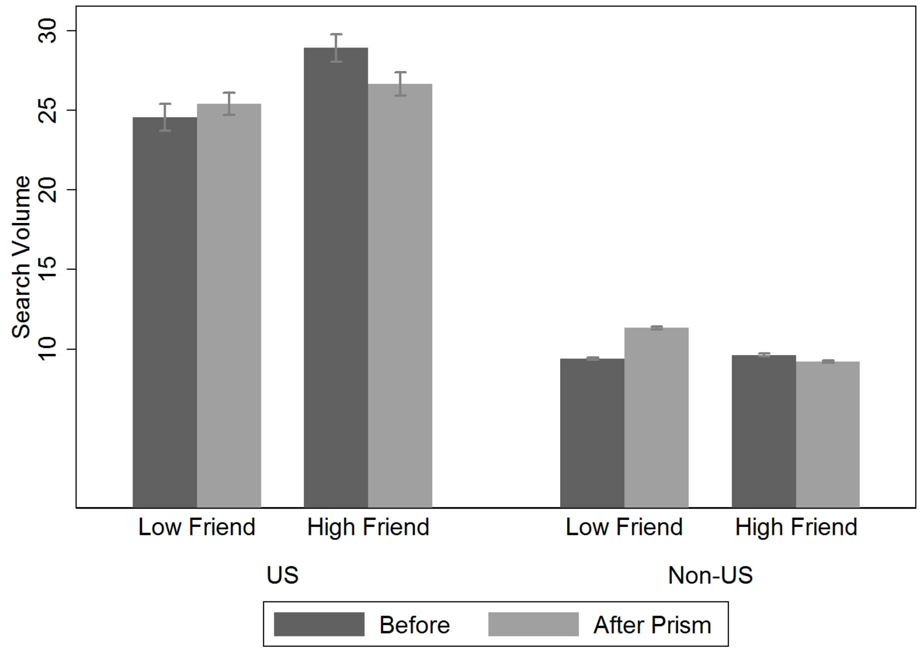


Figure 3: Search Volume Before and After PRISM Revelations

were rated as having a low level of likelihood that they would lead the user to be in trouble if a close friend knew about the user's search ("low-friend"), versus terms that had an above-median rating ("high-friend"). As shown by Figure 3, the overall pattern more or less holds: traffic for low-friend terms holds steady, and traffic for high-friend terms falls, though by less than in Figure 2 and in a less pronounced manner across the 40 non-US countries that we collected data for.

## 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:

$$\begin{aligned} SearchVolume_{ijt} = \beta TroubleCategory_i \times AfterPrism_t \\ +\gamma_i + \theta_j + \delta_t + \epsilon_i \end{aligned} \tag{1}$$

$\gamma$ is a series of fixed effects for each of our 245 keywords, $\theta_j$ is a series of fixed effects for each country, and $\delta_t$ is a series of fixed effects at the weekly level. The fixed effects $\gamma$ control for the different natural levels of search volume for each of the different search terms. $\theta_j$ captures general differences in search volume across countries. $\delta_t$ 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 $\beta$, which measures the differential in search volume for keywords that were more sensitive for this measure after the PRISM revelations. The main effect of $AfterPrism$ is collinear with the weekly fixed effects and is consequently dropped from our regressions. Similarly, our variable categorizing the search term according its 'trouble index', $TroubleCategory_i$, is collinear with the keyword fixed effects and is consequently dropped from the regression.

Table 6 presents our initial results for our weekly data over the period January 1, 2013 to December 31, 2013. The first three columns focus on a specification where we categorize our data based on median splits of the various trouble ratings.[18]

In particular, we isolate search terms which are above-median only in terms of 'getting you into trouble' with a friend (10% of the sample), search terms that are above-median only in terms of 'getting you into trouble with the government' (12% of the sample), and search terms which are above-median in terms of both forms of trouble (44% of the sample). Table 5 summarizes the average Likert-scale ratings for each of these categories and indicates that the above-median words on both scales were on average by far viewed as the most likely terms both for getting you into trouble with the US government and with a friend.

Table 5: 'Trouble' Rating of Google Search Terms by Trouble Categorization

| | No Trouble Mean | Gov Trouble Mean | Friend Trouble Mean | All Trouble Mean | Total Mean |
|---|---|---|---|---|---|
| Trouble Friend | 1.17 | 1.28 | 1.57 | 1.77 | 1.49 |
| Trouble Government | 1.20 | 1.65 | 1.29 | 1.93 | 1.58 |

Column (1) of Table 6 presents results for the US for the interaction between the indicator for the post PRISM revelations and these different classes of words. It suggests that there is a negative effect for the words that are perceived as having an above-median chance of getting you both into trouble with the US government and a friend. However, there is no negative effect for words which are perceived as troublesome in just a single dimension. This may be because of there being fewer words in these categories. However, it may also reflect the fact that, as shown in Table 5, the words that are above-median for both friend trouble and government trouble, are on average perceived as far more likely to provoke trouble with the

[18] In earlier versions of the paper we used the full indices rather than median splits, and obtained similar results.

US government. The point estimate suggests a decrease of approximately one index point from the baseline of 25 index points for these types of searches, or a four percent decrease in total for these words that are perceived as the most potentially troublesome. Overall, this provides 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 or with a friend.

Column (2) of Table 6 presents results where we demarcate the after PRISM period into both the first quarter and the second quarter after the PRISM revelations. This result suggests that the effect was the most pronounced in the first quarter after the PRISM revelations, but that it also persisted afterwards.

Column (3) of Table 6 examines whether there is any kind of pre-trend in the data looking at the previous month as an alternative 'Fake PRISM' start time. The coefficient for the placebo dummy for the month prior to the PRISM revelations is not significant. This again suggests there is no evidence of a measurable pre-trend in the data.

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 and a large variety of jurisdictions, so another news story relating to a small portion of them, such as an extreme weather event (for the DHS search terms) or a change in laws relating to childcare provision (for the local businesses terms) is unlikely to have shifted behavior for the whole. 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. Column (4) of Table 6 presents results where we just look a shorter ten-week window around the PRISM revelations. The estimate of the negative effect is slightly larger.

We also tried to rule out seasonality as being a driver of our results by repeating the

Table 6: In the US there was a decline in searches that were perceived as getting the searcher in trouble both with a friend and the US government

| | Base | Longer Period | Pre-Trend | Shorter Period 2013 | 2012 |
|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) |
| After Prism × Gov Trouble | 0.0863 | | -0.00408 | 0.391 | |
| | (0.290) | | (0.321) | (0.704) | |
| After Prism × Friend Trouble | 0.480 | | 0.518 | 0.128 | |
| | (0.354) | | (0.396) | (0.652) | |
| After Prism × All Trouble | -0.942** | | -0.943* | -1.460** | |
| | (0.313) | | (0.355) | (0.482) | |
| Quarter After Prism × Gov Trouble | | 0.118 | | | |
| | | (0.592) | | | |
| Quarter After Prism × Friend Trouble | | 1.376 | | | |
| | | (0.945) | | | |
| Quarter After Prism × All Trouble | | -1.374** | | | |
| | | (0.507) | | | |
| Second Quarter After Prism × Gov Trouble | | 0.0815 | | | |
| | | (0.294) | | | |
| Second Quarter After Prism × Friend Trouble | | 0.342 | | | |
| | | (0.339) | | | |
| Second Quarter After Prism × All Trouble | | -0.875** | | | |
| | | (0.319) | | | |
| After Fake Prism × Gov Trouble | | | -0.497 | | 0.492 |
| | | | (0.429) | | (0.498) |
| After Fake Prism × Friend Trouble | | | 0.208 | | 0.742 |
| | | | (0.536) | | (0.574) |
| After Fake Prism × All Trouble | | | -0.00887 | | 0.336 |
| | | | (0.463) | | (0.457) |
| Keyword Fixed Effects | Yes | Yes | Yes | Yes | Yes |
| Week Fixed Effects | Yes | Yes | Yes | Yes | Yes |
| Observations | 12740 | 12740 | 12740 | 2695 | 2695 |
| R-Squared | 0.925 | 0.925 | 0.925 | 0.958 | 0.967 |

OLS Estimates. Dependent Variable Is Search Volume Index As Reported By Google Trends. Weekly data over the period January 1, 2013 to December 31, 2013 in Columns (1)-(3). Weekly data for the ten week period around the revelations in Column (4). Weekly data for the same ten week period in 2012 in Column (5)

Robust Standard Errors Clustered At Search Term Level. $+\ p < 0.10, *p < 0.05, **p < 0.01, ***p < 0.001$. The main effects of $AfterPrism$ and the trouble category ratings are collinear with the week and keyword fixed effects and consequently both terms are dropped from the regression.

analysis of Column (4), using exactly the same June date but a year earlier in 2012. Column (5) of Table 6 repeats the analysis of Column (4) for 2012 to explore whether there are commonly such extreme drops, in these kind of searches perhaps as a result of seasonal variation. However, it finds no measurable effect. All the coefficients are reassuringly insignificant. This suggests that it is not seasonality brought about by comparing late spring with summer that is driving our results.

The results of Table 6, and in particular the fact that the negative effect of the PRISM revelations on searches is most pronounced around the time of revelations, raises the question of the extent to which this was driven simply by ongoing publicity surrounding the changes rather than a response to the information itself.

Table 7: Robustness Checks for the US Results to News Coverage

| | News Effects (1) | News+Short (2) | Log News+Short (3) |
|---|---|---|---|
| After Prism × Gov Trouble | -0.0322 | 0.433 | 2.272 |
| | (0.335) | (1.811) | (4.604) |
| After Prism × Friend Trouble | 0.123 | -1.160 | -6.291 |
| | (0.395) | (1.159) | (3.833) |
| After Prism × All Trouble | -0.690$^{+}$ | -2.352$^{*}$ | -5.474$^{+}$ |
| | (0.360) | (1.094) | (3.146) |
| Gov Trouble × Number of News Stories | 0.000189 | -0.0000316 | |
| | (0.000299) | (0.00101) | |
| Friend Trouble × Number of News Stories | 0.000569 | 0.000974 | |
| | (0.000483) | (0.000713) | |
| All Trouble × Number of News Stories | -0.000401 | 0.000674 | |
| | (0.000317) | (0.000624) | |
| Gov Trouble × Log News Stories | | | -0.252 |
| | | | (0.576) |
| Friend Trouble × Log News Stories | | | 0.861$^{+}$ |
| | | | (0.514) |
| All Trouble × Log News Stories | | | 0.538 |
| | | | (0.407) |
| Keyword Fixed Effects | Yes | Yes | Yes |
| Week Fixed Effects | Yes | Yes | Yes |
| Observations | 12740 | 2695 | 2695 |
| R-Squared | 0.925 | 0.958 | 0.958 |

OLS Estimates. Dependent Variable Is Search Volume Index As Reported By Google Trends. Weekly data over the period January 1, 2013 to December 31, 2013 in Column (1). Weekly data for ten-week window reported in Columns (2) and (3)
Robust Standard Errors Clustered At Search Term Level. $+\ p < 0.10, *p < 0.05, **\,p < 0.01, ***p < 0.001$.
The main effects of $AfterPrism$ and the trouble category ratings are collinear with the week and keyword fixed effects and consequently both terms are dropped from the regression.

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, though the introduction of these additional controls appear to introduce suggesting that the change we measure is not media-driven. Column (1) presents results for the full span of 2013. Column (2) presents results for the shorter window. Our results suggest that especially in the shorter period, the behavior we measure is not driven by news coverage.

Column (3) of Table 7 presents the results of using a logged measure to capture news coverage to potentially control for the role of extreme values. However, we caution that zeroes in our dataset are very prevalent.[19] 67.87% of our weeks did not have news stories concerning the PRISM revelations, so the log specification may be limited here. In general, news coverage seems to be negatively related to overall search volume, though none of our estimates are precisely estimated.

Another concern is whether the particular definition or the approach we took to the Mechanical Turk survey measures drove the results of Table 6. Table 8 investigates the robustness of our measures to different survey measures. One concern is that the categorization displayed in Table 5 into 'only friend trouble' 'only government trouble' and 'all trouble' drove the results. To investigate this, Column (1) of Table 8 presents the results of a simpler specification where we compare the results of an indicator for above-median ratings in the 'trouble with a friend category' and an indicator for above-median ratings in the 'trouble with the government' category, with no attempt to account between the potential overlap between the two. The results are similar to before, but we measure a negative effect for each indicator.

Column (2) of Table 8 investigates our results when we look at extreme values of the scale ratings: In this case, whether or not the rating was in the top decile. We observe a large and negative effect for the top decile of government trouble ratings, but do not observe an effect for the top decile of friend trouble ratings.

Another related concern is that our findings might be an artifact of the particular sensi-

[19] We dealt with this issue by simply adding 0.5 to all news metrics so that the log of news is measurable.

Table 8: Robustness Checks for the US Results (Alternative Definitions)

| | Two Categories (1) | Extreme Values (2) | Employer Trouble (3) | Family Trouble (4) |
|---|---|---|---|---|
| Post Prism × High Friend Trouble | -0.348$^{+}$ | | | |
| | (0.189) | | | |
| Post Prism × High Gov Trouble | -0.631$^{*}$ | | | |
| | (0.264) | | | |
| After Prism × Top Decile Friend | | 0.532 | | |
| | | (0.445) | | |
| After Prism × Top Decile Gov | | -1.316$^{*}$ | | |
| | | (0.535) | | |
| After Prism × Gov Trouble | | | -1.407$^{**}$ | |
| | | | (0.439) | |
| After Prism × Employer Trouble | | | -0.0117 | |
| | | | (0.392) | |
| After Prism × All Trouble | | | -0.688$^{*}$ | |
| | | | (0.290) | |
| After Prism × Gov Trouble | | | | -0.697$^{*}$ |
| | | | | (0.301) |
| After Prism × Family Trouble | | | | 0.598$^{+}$ |
| | | | | (0.319) |
| After Prism × All Trouble | | | | -0.689$^{*}$ |
| | | | | (0.322) |
| High Friend | 58.18$^{***}$ | | | |
| | (1.049) | | | |
| High Gov | -39.94$^{***}$ | | | |
| | (1.068) | | | |
| Keyword Fixed Effects | Yes | Yes | Yes | Yes |
| Week Fixed Effects | Yes | Yes | Yes | Yes |
| Observations | 12740 | 12740 | 12740 | 12740 |
| R-Squared | 0.925 | 0.925 | 0.925 | 0.925 |

OLS Estimates. Dependent Variable Is Search Volume Index As Reported By Google Trends.
Weekly data over the period January 1, 2013 to December 31, 2013
Robust Standard Errors Clustered At Search Term Level. $+\, p < 0.10, *p < 0.05, **\, p < 0.01 ***\, p < 0.001$.
The main effects of $AfterPrism$ and the trouble category ratings are collinear with the week and keyword fixed effects and consequently both terms are dropped from the regression.

tivity 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 the use, for example, of the term ‘friend’ that drove our results. Columns (3) and (4) of Table 8 investigates what happens when we use alternative measures of the ‘personal’ dimension of privacy, namely trouble with an employer and trouble with a family member. In both cases, we see a negative effect for the government trouble rating, and a somewhat smaller negative effect for words that were rated highly for both government trouble and trouble with the employer or family member. We speculate that the discrepancy in the results could be explained by Turkers rating 'time-wasting' searches highly in terms of likely employer trouble.[20]

One final concern is that Google could have strategically adjusted its algorithm so as to give less prominence to search results for a particular search term, in a manner that was contemporaneous with publicity about PRISM. This would affect clicks, but not the prior likelihood of a person entering a given search term. A deeper concern is that Google may have adjusted its search algorithm as a result of the search revelations, and in particular that the change in algorithm meant that people were more or less likely to search subsequently again for a different search term after the first set of search results failed to produce the results they were seeking. For example, it could be that the first search for 'pipe bomb' was rendered intentionally less informative, and so people searched again. To examine for this possibility, we went to Google Correlate, a Google database that allows a researcher to see what search terms are temporally correlated with a particular search term. We looked at the correlates of search terms for a random subsample of ten of the search terms in our data. The idea is that if the algorithm changed, we should see a difference in its accuracy, as reflected by how many times a Google users searches again as they were not able to find the result they were looking for. We could not see any pattern, however, that suggested a change in what search terms were used as substitutes for each other in June 2013, which would be suggestive of a change in the Google algorithm. As a final check, we also used a 'bounce back' metric, which measures whether or not a user searches again after performing

[20]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.

a Google search. We examined this using comScore data. However, we did not see after the revelations any change in the number of people going back to Google to search again for our search terms.

## 4 Moving to the International Context

Having established the robustness of our findings in the US, to understand both the mechanism and the broader context for our results we now turn to examine how our results change in the international context.

Table 9 provides an initial comparison of the US with its 40 top trading partners. Column (1) of Table 9 simply reports the results from Column (1) of Table 6 to provide a baseline for comparison. Column (2) of Table 9 reports the results for all other countries. Two patterns are clear in the data. First, the effect for words that might get you into trouble with the government without having a personal dimension of privacy-sensitivity, is now also negative and significant. Second, the point estimate for words that might get you into trouble with both a friend and the US government is less large than in the US estimates.

Columns (3) and (4) of Table 9 contrast our results between countries that speak English as their primary language outside of the United States and countries which do not. The negative effect of the PRISM revelations on searches is observable across all categories in English-speaking non-US countries. However, the effects for the non-English-speaking countries are confined to the categories where the words are perceived as likely to getting the person in trouble with the government. This suggests that familiarity with the language of their search terms and their nuances may drive the size of the chilling effect for personally-private search terms. It also suggests that outside of the US in countries which shared the English language, the chilling effects were more widespread across both government- and personally sensitive words compared to within the US.

Of course, one related concern may be that countries that are culturally removed from the US may have had a different pattern of chilling effects which we were not measuring because our Mechanical Turkers who rated the words were initially US-based. To explore this, we repeated our Mechanical Turk rating exercise with raters from India, and then used these ratings to explore the effect of the PRISM revelations ratings in India. The results are reported in Column (5) of Table 9. We observe a similar, albeit smaller, chilling effect to that observed in Column (1) for the US. This suggests that while of course relying on US measures of cultural sensitivities may miss some differences between cultures, country-specific measures of cultural sensitivity may not affect the major thrust of our results.

Table 9: Initial Results Across Trading Partners

| | US | Other Countries | English | Non English | India Ratings |
|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) |
| After Prism × Gov Trouble | 0.0863 | -0.513** | -0.640* | -0.488** | |
| | (0.290) | (0.160) | (0.283) | (0.155) | |
| After Prism × Friend Trouble | 0.480 | -0.0149 | -0.617* | 0.0457 | |
| | (0.354) | (0.0982) | (0.292) | (0.0933) | |
| After Prism × All Trouble | -0.942** | -0.215* | -0.641** | -0.201* | |
| | (0.313) | (0.0975) | (0.195) | (0.0960) | |
| After Prism × Gov Trouble | | | | | 0.0183 |
| | | | | | (0.200) |
| After Prism × Friend Trouble | | | | | -0.0459 |
| | | | | | (0.140) |
| After Prism × All Trouble | | | | | -0.422** |
| | | | | | (0.140) |
| Country Fixed Effects | No | Yes | Yes | Yes | No |
| Keyword Fixed Effects | Yes | Yes | Yes | Yes | Yes |
| Week Fixed Effects | Yes | Yes | Yes | Yes | Yes |
| Observations | 12740 | 509600 | 38220 | 484120 | 12740 |
| R-Squared | 0.925 | 0.634 | 0.837 | 0.637 | 0.943 |

OLS Estimates. Dependent Variable Is Search Volume Index As Reported By Google Trends over 2013.
Robust Standard Errors Clustered At Search Term Level. $+\ p < 0.10, *p < 0.05, **p < 0.01 ***p < 0.001$.
The main effects of $AfterPrism$ and the trouble category ratings are collinear with the week and keyword fixed effects and consequently both terms are dropped from the regression.

## 4.1 Familiarity with Surveillance

These consumer reactions may be a rational response to concerns about unwanted attention from the US government. However, the reaction of consumers in English-speaking non-US countries where search volume fell on personally sensitive search terms may also be appropriate, when placed in a broader and more behavioral context. It it obvious that consumers searching on "depression" in Canada are unlikely to get into direct trouble with the US government as a result of having conducted that search. However, increased knowledge of US government surveillance, for such users, may have had the effect of turning Google from a private space where it was possible to enter search terms freely, into a surveilled space where users are self-conscious of the fact that their searches may being monitored. Users are not used to policing their speech in a way that reflects a detailed knowledge of the search terms the US government is interested in. Lists of such search terms are usually not made public, and we may see from the DHS list that many of the search terms the US government is in fact interested in, do not appear to users to be sensitive in any respect. Therefore the response to surveillance among users that we document in all probability does not reflect detailed knowledge of the enormous volume of data that has been released on a dizzying array of different programs; but instead is more plausibly a generalized and intuitive response to the sense of being watched.

To examine how feelings about surveillance moderates the effect we find, we first turn to the 2014 edition 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?' and 'Is American monitoring of Your Country's leaders acceptable or unacceptable?' We then divided up the countries into whether they had above-median or below-median acceptance of US monitoring practices.

Table 10 investigates the difference between countries by their reported attitudes towards surveillance. Columns (1) and (2) compare a median split which distinguishes between the countries which were identified by the Pew survey as finding US government surveillance of foreign citizens unacceptable. It suggests that in the countries where US government surveillance was found to be less acceptable, there were larger effects. Columns (3) and (4) divide up the countries by whether they deemed surveillance of foreign leaders by the US government to be acceptable and again find more pronounced effects for countries which found surveillance less acceptable.

Table 10: Trading Partner Results by Acceptance of US Surveillance

| | Not Accept Monitor Citizens<br>(1) | Others<br>(2) | Not Accept Monitor Leaders<br>(3) | Others<br>(4) |
|---|---|---|---|---|
| After Prism × Gov Trouble | -0.572** | -0.429** | -0.568** | -0.426** |
| | (0.187) | (0.141) | (0.187) | (0.139) |
| After Prism × Friend Trouble | -0.110 | 0.0996 | -0.103 | 0.103 |
| | (0.113) | (0.105) | (0.112) | (0.107) |
| After Prism × All Trouble | -0.344** | -0.127 | -0.330** | -0.131 |
| | (0.116) | (0.0921) | (0.114) | (0.0921) |
| Country Fixed Effects | Yes | Yes | Yes | Yes |
| Keyword Fixed Effects | Yes | Yes | Yes | Yes |
| Week Fixed Effects | Yes | Yes | Yes | Yes |
| Observations | 254800 | 267540 | 267540 | 254800 |
| R-Squared | 0.661 | 0.616 | 0.658 | 0.615 |

OLS Estimates. Dependent Variable Is Search Volume Index As Reported By Google Trends over 2013. Robust Standard Errors Clustered At Search Term Level. $+\ p < 0.10, *p < 0.05, **p < 0.01, ***p < 0.001$. The main effects of $AfterPrism$ and the trouble category ratings are collinear with the week and keyword fixed effects and consequently both terms are dropped from the regression.

This suggests that the results we find are driven by the level of acceptance of surveillance. One other way to approach this is to compare countries where own-government surveillance is commonplace with those where it is not. This led to our second approach, which was to obtain data on how habituated a country's citizens are to surveillance.[21]

Table 11 reports in Columns (1) and (2) the results of dividing our country data into countries which mandate that their adult citizens have to possess an ID, and in Columns (3) and (4) the results of dividing our country data into countries which mandate that their adult citizens have to carry their ID. These results suggest an interesting nuance, which is that it is users in the countries which are not used to surveillance from their own government in terms of a national ID, for whom there was a more profoundly negative effect, especially

[21]We do not intend to imply by this that countries whose citizens find mass surveillance more abhorrent, are coterminous with countries that do not in fact conduct mass surveillance; it may also be the case that countries whose citizens do not find mass surveillance so abhorrent are countries whose governments are more open about the within-country surveillance they conduct.

for the personally sensitive terms, more than for users in countries which require citizens to own or carry a national ID.

Table 11: Trading Partner Results by Familiarity with Surveillance

| | Mandatory ID: Yes | No | Mandatory Carry ID: Yes | No |
|---|---|---|---|---|
| | (1) | (2) | (3) | (4) |
| After Prism × Gov Trouble | -0.425** | -0.627** | -0.628** | -0.485** |
| | (0.144) | (0.198) | (0.181) | (0.161) |
| After Prism × Friend Trouble | 0.129 | -0.232 | 0.198 | -0.0246 |
| | (0.0861) | (0.151) | (0.156) | (0.0981) |
| After Prism × All Trouble | -0.124 | -0.422** | 0.0700 | -0.266* |
| | (0.0898) | (0.126) | (0.113) | (0.100) |
| Country Fixed Effects | Yes | Yes | Yes | Yes |
| Keyword Fixed Effects | Yes | Yes | Yes | Yes |
| Week Fixed Effects | Yes | Yes | Yes | Yes |
| Observations | 331240 | 191100 | 50960 | 471380 |
| R-Squared | 0.638 | 0.648 | 0.722 | 0.625 |

OLS Estimates. Dependent Variable Is Search Volume Index As Reported By Google Trends over 2013. Robust Standard Errors Clustered At Search Term Level. $+\ p < 0.10, *p < 0.05, **\ p < 0.01 ***\ p < 0.001$. The main effects of $AfterPrism$ and the trouble category ratings are collinear with the week and keyword fixed effects and consequently both terms are dropped from the regression.

## 4.2 Diplomatic Closeness to the US

One important question for interpreting the results of Table 9 which suggest that internationally citizens were less likely to search on search terms perceived as likely to get them into trouble with the US government, is the extent to which they reflect a suppression in searches on problematic search terms by actual potential antagonists to US interests - for example searching for information about a pipe bomb as part of a violent attack - or instead a suppression of searches by citizens of those countries who are merely interested in finding out more about the topics of these search terms. To investigate this further, we study whether our search terms vary by the extent to which the country is considered an antagonist of the US.

To find a potential source for lists of countries that might be considered antagonistic to US interests, we searched for an external metric for the countries that the NSA considers to be priorities for surveillance. We found such a listing in a document called the 'National Intelligence Priorities Framework.' For understandable reasons, the government does not publicize its priority targets, but enough information has been leaked over time for an anonymous hacker to have collated the available information into a viable data source.[22] We used data from this complex framework to construct a five-point scale ranking countries by priority for US intelligence collection. The countries included in our study with a 'one,' or top, ranking are Brazil, China, India, Iran, Mexico, Pakistan and Russia.[23]

There are some countries which might be regarded as having many Google users antagonistic to US interests, like Afghanistan and Iraq, which are are not part of our dataset, as their countries are not among the US's top 40 trading partners. However, the characteristics which are endogenous to users in a country being presumed to be more antagonistic to US interests, would in turn affect the volume of Google searches being conducted in the country. In several of these countries there is aerial bombing by the US and other insurgent action, which probably distorts and reduces Google searches conducted by those countries' citizens. In North Korea, its isolationist strategy means that there are only 28 websites in total[24] and there is no official access to Google at all. However, this general issue of interpretation - that is, that the antagonist country's response to the revelations would be naturally suppressed because of endogenous characteristics that are correlated with its being an antagonist country - is a limitation which should be recognized when interpreting our results.

Table 12 reports the results of dividing the 40 trading partners by their relative importance for US intelligence. It is clear that the smallest effect of the surveillance revelations

[22] See `https://cryptome.org/2013/11/nsa-nipf-v3.htm` for details.

[23] Countries with a 'two' ranking are Israel, Saudi Arabia, South Korea, Taiwan, and Turkey. Countries with a 'three' ranking are France, Germany and Japan. Colombia, Egypt, Italy, Malaysia, Nigeria, Singapore, Spain, Sweden and Thailand have a 'four' ranking; Denmark, Norway and the Philippines have a 'five' ranking; and the data source, by omission, implies that the remaining countries in our study are such low priorities for intelligence collection that they are not ranked at all.

[24] McGoogan, C.,'North Korea Revealed To Have Just 28 Websites', The Telegraph, September 21, 2016, available at `http://www.telegraph.co.uk/technology/2016/09/21/north-koreas-internet-revealed-to-have-just-28-websites/`, accessed on February 16, 2017.

Table 12: Trading Partner Results by Intelligence Priorities Ranking

| | Rank 1 | Rank 2 | Rank 3 | Rank 4+ |
|---|---|---|---|---|
| | (1) | (2) | (3) | (4) |
| After Prism × Gov Trouble | -0.287* | -0.286+ | -0.530* | -0.593** |
| | (0.117) | (0.147) | (0.198) | (0.181) |
| After Prism × Friend Trouble | 0.169+ | 0.149 | -0.463* | -0.0253 |
| | (0.0929) | (0.102) | (0.175) | (0.118) |
| After Prism × All Trouble | -0.0848 | -0.274* | -0.284* | -0.259* |
| | (0.0694) | (0.105) | (0.139) | (0.112) |
| Country Fixed Effects | Yes | Yes | Yes | Yes |
| Keyword Fixed Effects | Yes | Yes | Yes | Yes |
| Week Fixed Effects | Yes | Yes | Yes | Yes |
| Observations | 89180 | 63700 | 38220 | 331240 |
| R-Squared | 0.580 | 0.718 | 0.709 | 0.656 |

OLS Estimates. Dependent Variable Is Search Volume Index As Reported By Google Trends over 2013. Robust Standard Errors Clustered At Search Term Level. $+\ p < 0.10, *p < 0.05, **p < 0.01, ***p < 0.001$. The main effects of $AfterPrism$ and the trouble category ratings are collinear with the week and keyword fixed effects and consequently both terms are dropped from the regression.

is found for the countries that ranked as highest surveillance priorities. By contrast, the largest effect of the surveillance revelations is found for countries that are low surveillance priorities and traditional allies of the US. One interpretation of these results is that users in the countries that were the highest intelligence priorities of the US already believed they were being monitored, so the revelations had little effect on behavior. In any case, these results suggest that in all likelihood the 'chilling effects' that we document were not the result of the suppression of searches by potential bad actors, but instead the result of citizens self-censoring more innocently intentioned queries, after receiving information that search queries were being monitored.

To summarize, the use of international data allows a more nuanced understanding of what drives the effects that we measure. We show that in the international context that the effects are stronger in regions which share a first language of English with, and therefore are closer in culture to, the US. We also show that in these countries, the effect is uniformly negative for words that are considered sensitive both in the personal and government domain. One potential explanation for why the revelations had a smaller effect in culturally dissimilar countries to the US is their being more accustomed to surveillance.

We find that our effects are larger and more significant when we compare countries that both express the idea that US surveillance of their country is ‘unacceptable’ and also are less used to their own governments performing surveillance on them. We then turn to try and understand whether the presence of a chilling effect should be thought of as simply a rational response of potential antagonists to US interests who were worried about being watched, or instead was a less rational response of consumers who were not potential antagonists but instead were uncomfortable with being monitored by the US government. We found that, perhaps counter-intuitively, the biggest negative effects were not in countries which are US government intelligence targets due to potential concerns about their threat level to US interests. Instead, the biggest chilling effects were found in countries which are traditionally considered US allies. This may reflect that in countries which are considered higher threats to US interests, citizens could already have been cautious about their use of a US search engine, whereas citizens in countries which are traditionally allies of the US may not have known about the potential for US government surveillance, and their response would therefore reflect an instinctive unease at being watched or potentially getting into trouble for an innocently intended search query.

# 5 Surveillance and the International Competitiveness of Search Engines

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. In this section, we explore the potential economic ramifications of the effect we find.

## 5.1 Did Users Switch Search Engines?

One natural question for understanding the effect is the extent to which this should be regarded as reflecting a chilling effect on general searching behavior, or whether it instead reflects the potential for users to switch to other search engines which are perceived as not being subject to surveillance by the US government. Of course in the US, all major search engines were being monitored - a fact that was clear as part of the PRISM revelations. However, outside of the US there were search engines available to citizens of other countries that were not part of the PRISM program. We explore the effect of such outside options in this section.

Table 13: Trading Partner Results by Prevalence of External Search Engines

| | Google First | Google Not First | No Other Top 5 Search | Other Top 5 Search |
|---|---|---|---|---|
| | (1) | (2) | (3) | (4) |
| After Prism × Gov Trouble | -0.515** | -0.299*** | -0.600** | -0.402** |
| | (0.170) | (0.0825) | (0.185) | (0.144) |
| After Prism × Friend Trouble | -0.00489 | 0.0230 | -0.102 | 0.0915 |
| | (0.107) | (0.0785) | (0.133) | (0.0833) |
| After Prism × All Trouble | -0.243* | -0.104 | -0.336** | -0.134 |
| | (0.105) | (0.0648) | (0.118) | (0.0873) |
| Country Fixed Effects | Yes | Yes | Yes | Yes |
| Keyword Fixed Effects | Yes | Yes | Yes | Yes |
| Week Fixed Effects | Yes | Yes | Yes | Yes |
| Observations | 484120 | 38220 | 254800 | 267540 |
| R-Squared | 0.639 | 0.634 | 0.659 | 0.618 |

OLS Estimates. Dependent Variable Is Search Volume Index As Reported By Google Trends over 2013. Robust Standard Errors Clustered At Search Term Level. $+\ p < 0.10, *p < 0.05, **p < 0.01 ***p < 0.001$. The main effects of $AfterPrism$ and the trouble category ratings are collinear with the week and keyword fixed effects and consequently both terms are dropped from the regression.

In Table 13 we explore how differences in the outside availability of search engines affects the results. 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. 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 Russia, there is also a national competitor named Yandex. Columns (1) and (2) contrast searching behavior before and after the PRISM revelations in countries where Google is ranked as the number one search engine by Alexa and in countries where Google is not ranked as the number one search engine. Columns (3) and (4) compare searching behavior where there is another search engine in the top five most popular websites as ranked by Alexa for that country. Both comparisons suggest that the effects we measure are stronger than when there is less of an obvious outside option to Google. In other words, it seems that users are more likely to self-censor their behavior on Google when there they have fewer easy alternatives to Google.

Table 13 suggests that there is less of an effect for countries such as China where Google is not heavily used. However, we caution that there may be many explanations for the direction of this effect, due to the endogenous nature of Google's presence or absence in a country. For example, the reason that that Google does not have a sizable presence in China is that the Chinese government restricts internet access, which is in turn related to their desire to monitor their citizens. Similar concerns may also be present in Russia - another country where there is both customary government surveillance and an alternative search engine to Google.

It is also natural to ask in the US whether, because we simply look at changes within Google, we miss changes that may have occurred in terms of people switching from Google to other search engines.

This is a concern because 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.[25]

[25]See "Anonymous search tool DuckDuckGo answered 1bn queries in 2013", The Guardian, January 9, 2014, available at `https://www.theguardian.com/technology/2014/jan/09/anonymous-search-tool-duckduckgo-1bn-queries-2013-google`, accessed February 16, 2017.

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 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.[26] Since Tor itself may successfully mask country of origin data, there is a small chance of an increase in measurement error if there were an increase in Tor usage after the Snowden revelations, but that measurement error would work against our ability to estimate an effect.[27]

Another way that users may potentially try and protect their privacy online is to use a private browsing mode. However, as pointed out by (Aggarwal et al., 2010), these private browsing modes do not actually prevent a website such as Google from tracking what you search for and what country you are from.

## 5.2 Were the Search Terms Affected Economically Meaningful?

Personally-sensitive search queries may be more valuable to advertisers, as consumers may use search engines to research these matters rather than consulting friends or family or other public resources, giving advertisers few channels other than search engine advertising to reach consumers. Some of the health-related search terms we study are profitable keywords for search engines.[29] We studied whether the after PRISM revelations effect we measure varied with whether or not the search term was health-related. Twenty percent of the search terms in our database were health-related. Just under half of these terms fell in the highest category of perceived trouble, which is high trouble with both the government and a friend. We investigate the difference in measured effect between health-related search terms and

[26]See `https://metrics.torproject.org/` for Tor usage statistics for the relevant timeframe.

[27]The overall proportion of Web traffic that is encrypted has also risen sharply from around 10% at the time of the PRISM revelations to around 70% of web traffic at the end of 2015.[28], but that rise does not affect Google Trends' data.

[29]See "How Does Google Make Its Money: The 20 Most Expensive Keywords in Google AdWords", available at `http://www.wordstream.com/articles/most-expensive-keywords`, accessed February 17, 2017. This report suggests that four of the top twenty categories reflected the seeking of medical treatment.

other search terms in Table 14. We were not able to use the full trouble categorization depicted in Table 5, because we did not have sufficient variation in terms of health search terms that were considered likely to get you into trouble with the US government (and not likely to get you into trouble with a friend). What we observe is that the negative effect is larger both in the US and elsewhere for these health-related terms, and in the US is considerably larger. This of course suggests that the chilling effect that we document may have been larger for some of the search terms which were most profitable for search engines to monetize. It also has potentially larger implications for the economic well-being of a country's citizens, if they were deterred from seeking out health information online.

Table 14: Responses to PRISM Revelations are Larger for Health Related Search Terms

| | US:Non-Health | US:Health | Other:Non-Health | Other:Health |
|---|---|---|---|---|
| | (1) | (2) | (3) | (4) |
| After Prism × Both High Trouble | -0.718$^{+}$ | -2.302$^{***}$ | -0.0552 | -0.291$^{***}$ |
| | (0.365) | (0.424) | (0.105) | (0.0572) |
| Country Fixed Effects | No | No | Yes | Yes |
| Keyword Fixed Effects | Yes | Yes | Yes | Yes |
| Week Fixed Effects | Yes | Yes | Yes | Yes |
| Observations | 10140 | 2600 | 405600 | 104000 |
| R-Squared | 0.908 | 0.971 | 0.645 | 0.615 |

OLS Estimates. Dependent Variable Is Search Volume Index As Reported By Google Trends over 2013. Robust Standard Errors Clustered At Search Term Level. $+\ p < 0.10, *p < 0.05, **p < 0.01, ***p < 0.001$. The main effects of $AfterPrism$ and the trouble category ratings are collinear with the week and keyword fixed effects and consequently both terms are dropped from the regression.

## 5.3 Absolute Magnitudes of the Effects

One drawback of the focus in our empirical analysis on the Google Trends data is that it is an index, which does not allow clear conclusions to be made about the absolute size of the effects that we are measuring. To address this, we gathered data for the two months surrounding the Snowden revelations from comScore's Search Planner database.[30] This database reports

[30] ComScore tracks the online activity of a panel of more than two million users based in the US and subsequently aggregates their search patterns to the search-term level for resale to commercial clients. ComScore recruits its panel members through affiliate programs and partnering with third party application providers.

the average click behavior of consumers following a keyword search on major search engines. For each keyword search, comScore reports the monthly aggregate number of clicks received by a selection of websites.

This data has the advantage that it reports absolute numbers of clicks performed after a Google search. However, it has the disadvantage that because it is focused on the US, we cannot use it to measure the international effects which are the focus of our study. That said, we did recreate our effects for the US in Figures A-1 and A-2, which are reported in the appendix. In the raw data, average clicks on search terms that were likely to lead to trouble with the US government fell from 931,944 clicks to 808,544 clicks after the revelations. We emphasize, though, as is clear from the figures, that the aggregated and limited nature of the comScore data precludes any claim of statistical significance for this drop.

These are average figures, but there are many searches in our data which are classified as more sensitive and attract a relatively large number of searches. According to comScore, there are multiple terms that received over one million clicks in the US as a result of consumers searching for them which also have above-median ratings in terms of likelihood to get you into 'trouble with a friend' personal privacy metric. These included bankruptcy, burn, chemical, cutting, depression, gas, guns, infection, nuclear, pregnant, revolution, shooting, suicide, weed and weight loss.

The other advantage of collecting this data from comScore is unlike Google Trends data, they collect data from other search engines. This allowed us to do two things. First, it allowed us to see whether the effect extended to other search engines. Figures A-3 and A-4 in the appendix suggest that at least in the raw data, the direction of the change is similar on Yahoo, AOL and Bing, though the sparseness of the data prevents any conclusions about statistical significance.

---

ComScore emphasizes and discusses the representativeness of their sample to the general population in their Marketer User Guide. The comScore data has also been used in several academic studies (Montgomery et al., 2004; Gentzkow and Shapiro, 2011; Santos et al., 2012).

There is one search engine in the comScore data that was not explicitly mentioned in the PRISM slide, which is Ask.com. Ask.com was previously known as AskJeeves.com and also operates via the Ask.com toolbar, and has the smallest number of users of the search engines tracked by comScore. We looked to see whether we could observe an change in search behavior on the Ask platform. As shown by Figures A-5 and A-6 in the appendix, there was no change in volume of search queries based on the degree of privacy sensitivity that was observable in the data. Of course, we do caution that this result is based on sparse data: A likely reason for Ask.com's search engine's non-inclusion in PRISM is its small size. However, it is at least reassuring as to the mechanism that we do not observe the same change as we observe for the Google search engine in the raw data.

This section has documented a variety of economic implications of the chilling effects that we study. The first implication is that these chilling effects are not driven by the presence of an easy outside option in the form of an nationally-based alternative to Google. Therefore, rather than the reduction in volume of search engine searches representing switching behavior, it seems more likely they represent self-censorship and a decision to not actually search on a term. Second, we show that much of the negative effect for the words which were considered troublesome in both the government and the private domain appears to be driven by health-related search terms. This has economic implications both for Google, because health-related keywords tend to be profitable to monetize, but also has broader implications for economic welfare if citizens do not seek out health information. Third, we show some evidence, which, albeit imprecisely, suggests that what we measure had real absolute consequences in terms of lower absolute volumes rather than simply changes in relative volumes of searches.

# 6 Conclusion

This study is the first to provide substantial empirical documentation of a chilling effect, both domestically and internationally, that appears to be related to increased awareness of government surveillance.[31] Furthermore, this chilling effect appears to apply to search behavior that is not strictly related to the government but instead forms part of the private domain among English-speaking non-US countries, and is driven in part by a reduction in searches that are related to health.

These results should be set in the broader context of other evidence of the effects for US international competitiveness of NSA surveillance. For example, the Indian government ruled out a partnership with Google on the basis of NSA surveillance.[32] Similarly, there has been increased uncertainty regarding the transfer of data by US companies between the US and Europe due to the invalidation of the 'Safe Harbor' agreement as a result of a European Court of Justice ruling based on the revelations; the Safe Harbor agreement previously allowed US companies to self-certify that they were broadly complying with data protection principles with respect to the data they held on EU citizen users. Our results suggest that in addition to these more visible manifestations of the effects of surveillance on competitiveness, policy makers should also account for the potential for less visible effects on consumer behavior when evaluating the desirability of mass surveillance.

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, our results are focused on

[31] Such effects may not be limited simply to Google's or other companies' 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 (Dinev et al., 2008), no systematic study yet exists.

[32] Tripathy, D. and Gottipati, S., "India's election regulator drops plan to partner Google", Reuters, January 10, 2014, available at `http://www.reuters.com/article/us-india-elections-google-idUSBREA080YE20140110`, accessed February 17, 2017.

the effects of revelations about government surveillance as opposed to the direct effects of government surveillance *per se*. Third, though we show the effects for a limited subset of search queries, we do not know how extensive the effects are across the universe of searches. 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. *Mimeo, CMU*.

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*.

Aggarwal, G., E. Bursztein, C. Jackson, and D. Boneh (2010). An analysis of private browsing modes in modern browsers. In *USENIX Security Symposium*, pp. 79–94.

Athey, S., C. Catalini, and C. Tucker (2017). The digital privacy paradox: Small money, small costs, small talk. *Mimeo, MIT*.

Barnes, S. B. (2006). A privacy paradox: Social networking in the United States. *First Monday 11*(9).

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., A. Goldfarb, and C. Tucker (2015). Privacy regulation and market structure. *Journal of Economics & Management Strategy 24*(1), 47–73.

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.

Cohen, R. (2013, July). NSA is doing what Google does. *Washington Post*.

Cooper, J. C. (2017). Anonymity, Autonomy, and the Collection of Personal Data: Measuring the Privacy Impact of Google's 2012 Privacy Policy Change. *Mimeo, George Mason University*.

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. *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*.

Gentzkow, M. and J. M. Shapiro (2011). Ideological segregation online and offline. *The Quarterly Journal of Economics*.

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, January). Privacy regulation and online advertising. *Management Science 57*(1), 57–71.

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: 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. R. and C. Tucker (2009, July). Privacy Protection and Technology Diffusion: The Case of Electronic Medical Records. *Management Science 55*(7), 1077–1093.

Montgomery, A., S. Li, K. Srinivasan, and J. Liechty (2004). Modeling online browsing and path analysis using clickstream data. *Marketing Science 23*(4), 579–595.

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

Penney, J. (2016). Chilling effects: Online surveillance and Wikipedia use. *Berkeley Technology Law Journal 31*(1), 117.

Richards, N. (2013, May). The dangers of surveillance. *Harvard Law Review 1934*, 126.

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

Santos, B. D. L., A. Hortacsu, and M. Wildenbeest (2012). Testing models of consumer search using data on web browsing and purchasing behavior. *American Economics Review 102*(6), 2955–80.

Somin, I. (2016). *Democracy and political ignorance: Why smaller government is smarter*. Stanford University Press.

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

# A Appendix

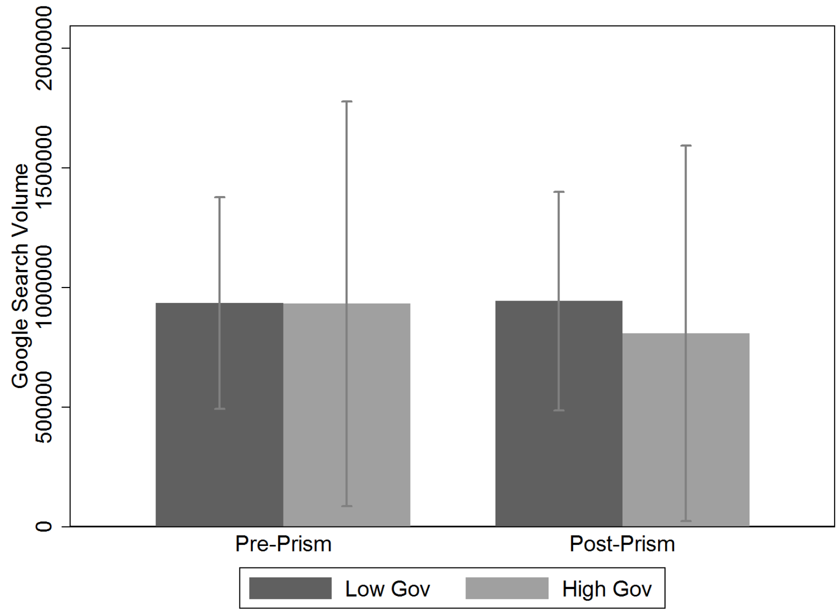


Figure A-1: Google Search Engine: US change in Search Volume by Likelihood of Trouble with the US Government from May to June 2013 - data from comScore

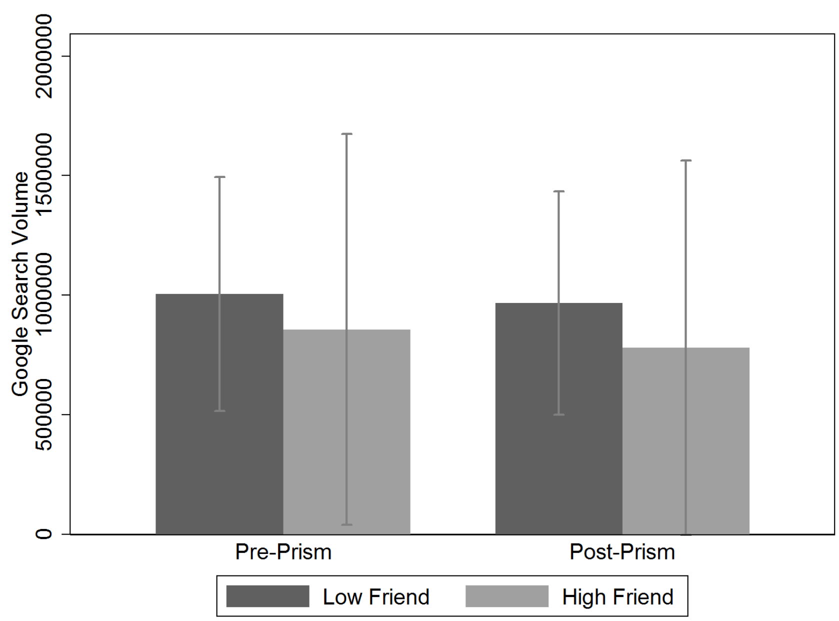


Figure A-2: Google Search Engine: US change in Search Volume by Likelihood of Trouble with a Friend from May to June 2013 - data from comScore

Table A-1: 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 A-2: 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 A-3: 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 A-4: 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 A-5: Google Search Terms

| | Friend Trouble Rating |
|---|---|
| arcade | 1 |
| beautysalon | 1.22 |
| butcher | 1.22 |
| childcare | 1 |
| cleaners | 1 |
| gardener | 1 |
| gym | 1 |
| movies | 1 |
| nursinghome | 1 |
| restaurant | 1 |
| thriftstore | 1 |
| weather | 1 |
| Total | 1.04 |

Table A-6: Cross-correlation table

| Variables | Trouble Empl | Trouble Family | Trouble Friend | Trouble Gov |
|---|---|---|---|---|
| Trouble Empl | 1.00 | | | |
| Trouble Family | 0.86 | 1.00 | | |
| Trouble Friend | 0.84 | 0.94 | 1.00 | |
| Trouble Gov | 0.79 | 0.65 | 0.70 | 1.00 |

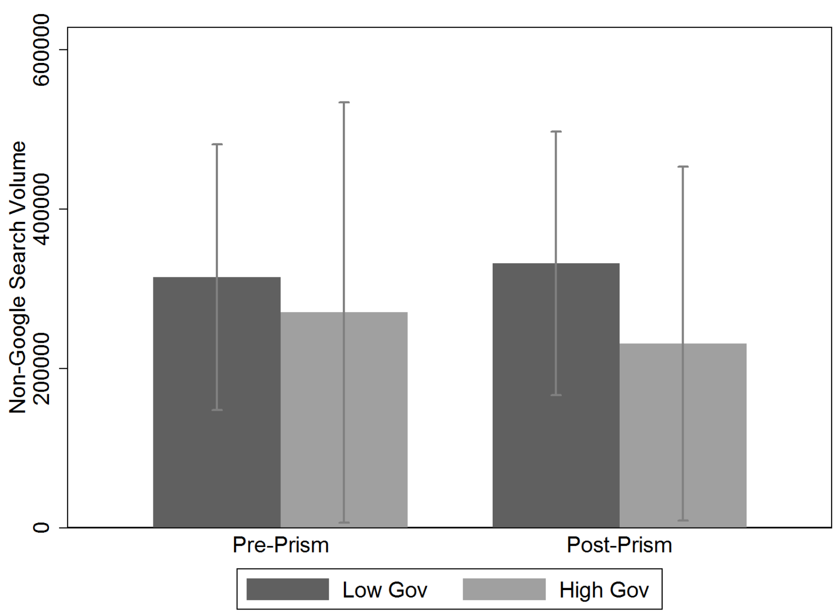


Figure A-3: Non-Google Search Engines: US change in Search Volume by Likelihood of Trouble with the US Government from May to June 2013 - data from comScore

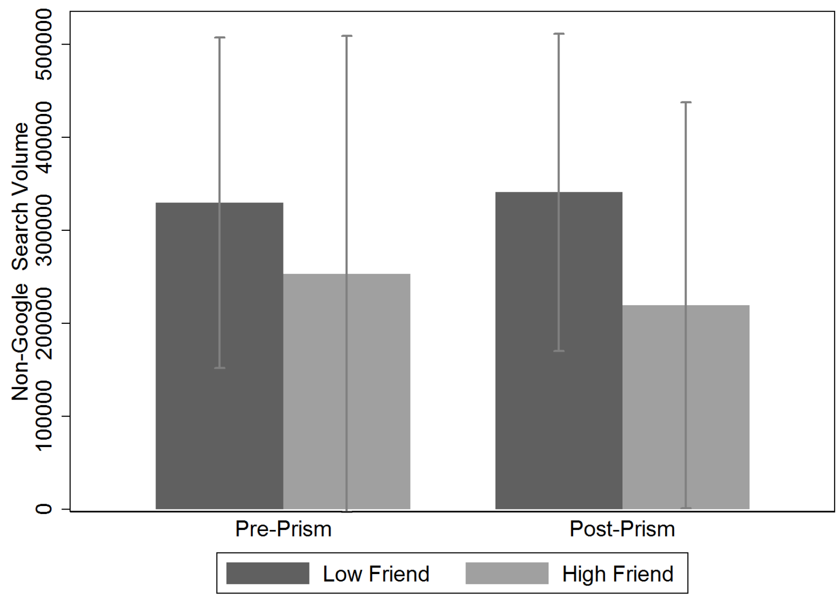


Figure A-4: Non-Google Search Engines: US change in Search Volume by Likelihood of Trouble with a Friend from May to June 2013 - data from comScore

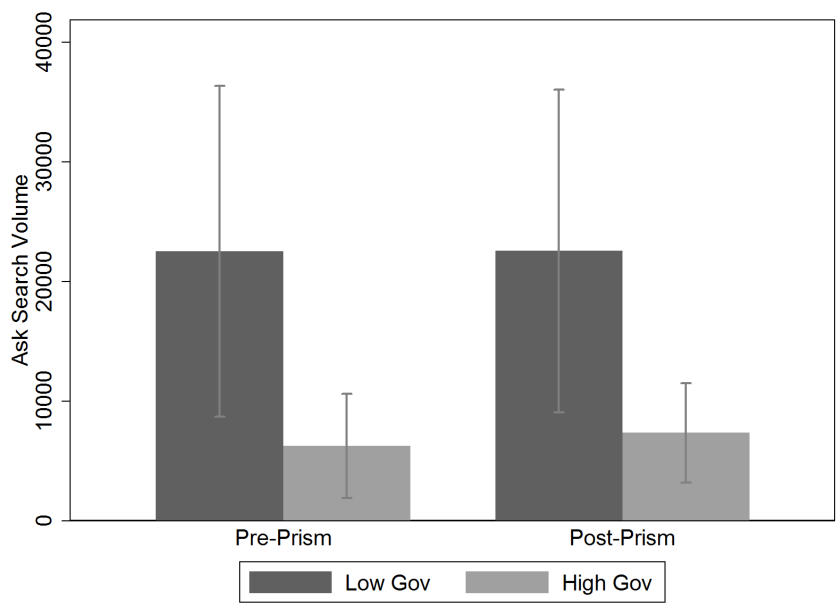


Figure A-5: Ask.com Search Engine: US change in Search Volume by Likelihood of Trouble with the US Government from May to June 2013 - data from comScore

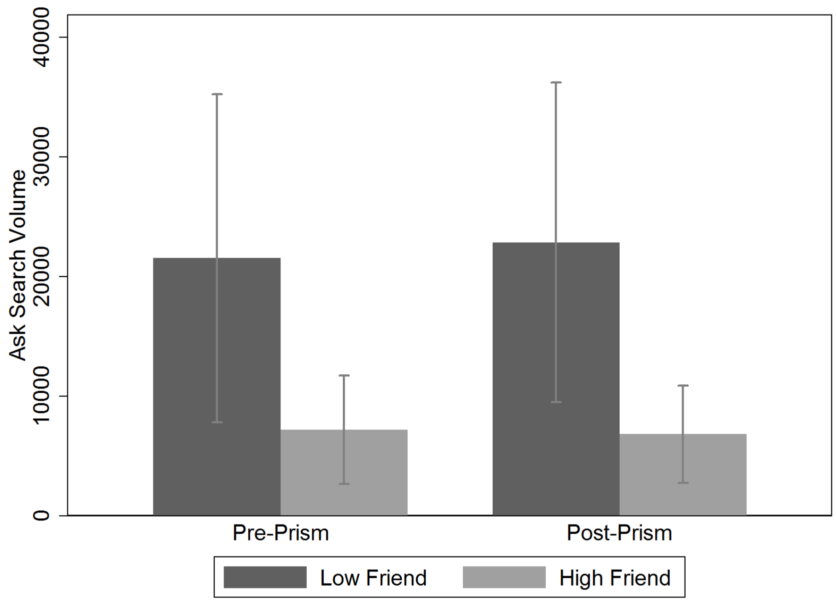


Figure A-6: Ask.com Search Engine: US change in Search Volume by Likelihood of Trouble with a Friend from May to June 2013 - data from comScore