

An Empirical Analysis on the Use and Reporting of National Security Letters

Alex Bellon, Miro Haller, Andrey Labunets, Enze Liu, and Stefan Savage¹

UC San Diego

{abellon,mhaller,alabunets,e7liu,savage}@ucsd.edu

¹Authors listed alphabetically.

Abstract

National Security Letters (NSLs) are a form of legal process that empowers parts of the United States federal government to request certain information for national security purposes. Authorized under five distinct statutory provisions, NSLs are similar to administrative subpoenas and can be issued directly by elements of the executive branch without requiring prior approval from a court or grand jury. Importantly, NSLs authorize the imposition of nondisclosure orders (aka “gag orders”) on the receiving party. Controversy about potential abuses of this authority has driven a range of legal and policy discussions. To address these concerns, both the public sector (through Congressional reporting requirements) and the private sector (through corporate “transparency reports”) have sought to document the usage of NSLs in aggregated form. In addition, *some* NSLs are now made public due to changes in the law that placed conditions on the nondisclosure requirements of NSLs. However, each of these disclosures is limited in scope, time, and kind. It remains unclear to what extent the heterogeneous data from different sources can be combined to draw meaningful conclusions about the usage of NSLs.

In this paper, we attempt to improve this state of affairs by consolidating the available data around NSLs from various sources and analyzing it. Our data collection allows us to answer two questions: (1) what can the public effectively learn from the reported data and does this information suffice to assess the NSL usage? and (2) how accessible is this data collection? On the one hand, we show that *after* consolidating and processing data, longitudinal trends in the usage of NSLs can be observed. For instance, we find a significant increase in NSL requests for non-US persons and that the policy reforms to decrease the mandated nondisclosure period appear to be effective. The observed trends suggest that the current transparency mechanisms are viable safeguards against the excessive use of NSLs. On the other hand, we find that aggregating and normalizing the data is a challenging task, and one of our main contributions is producing a normalized and accessible data set and making it available for future research. The various data sources pose major challenges as they all lack a standardized and machine-readable format. Releasing transparency reports and NSLs, after their gag order was lifted, is optional for companies and only done consistently by a few. While the government data is comparatively straightforward to find, as the publishing entities are known, processing them is challenging. Indeed, the format of the reported data changed repeatedly and the publication as scanned PDFs with continuous text requires manual reviewing, parsing, and validating of the data. We even find inconsistencies within and across data sources. Overall, the laborious data collection process hinders external and internal auditing efforts and demonstrates the need for a unified and more usable dataset for NSLs. We advocate for publishing NSL statistics in a standardized and easily processable format to facilitate third-party data analysis that ensures that Congress and the public can satisfy their role of overseeing the use of NSLs by government agencies.

1 Introduction

National Security Letters, or NSLs, are legal requests issued by US government agencies (predominantly the FBI) to acquire information for investigating national security issues. As statutorily authorized, named agencies can request information about individuals from a broad array of third-party sectors, including financial institutions, communication providers, and credit agencies without requiring approval from the courts or grand jury.

As originally constituted, NSLs were a powerful investigative tool with little judicial oversight, relaxed reporting requirements, and nondisclosure provisions preventing companies from reporting the reception of NSLs or their content. Perhaps due to these properties, NSL usage grew rapidly [26], and lower courts observed that the NSL practices might conflict with the First Amendment and the separation of powers [43, 42]. Subsequently, Congress attempted to increase transparency and congressional oversight for NSLs, starting with the USA Patriot Improvement and Reauthorization Act statutes in 2006. After Edward Snowden’s disclosures concerning government surveillance activities in 2013 [29], companies further pushed back against NSL, fighting for more transparency, and obtained permission to publish their own NSL statistics (albeit quantized in buckets of size 250 or more to obscure precise analysis) [20, 15].

While there long have been efforts to create a more modern and transparent government, these developments have not yet been adopted for NSL statistics. Starting in 2009, the Open Government Initiative aims to “[i]mprove public access to government data, research, and information to enhance transparency, accountability, and equitable outcomes” [30] and Data.gov [25] hosts over 290,000 raw data sets. In comparison, the NSL data does not follow any of the best practices for publishing structured, machine-readable government data discussed by Joe Calandrino and Harlan Yu in a series of blog posts [23]. The NSL information is published by different government bodies, in PDFs that require manual labor to parse, and with frequently changing data formats. This lack of a universal data format creates an obstacle to analysis, as it requires a significant amount of labor to acquire, extract, and normalize the information scattered across different reports. Perhaps, security concerns disincentivize agencies from releasing too fine-grained NSL statistics. However, while the current unstructured data publishing does not preclude analysis, the intricate and cumbersome process of doing so hinders the public from fulfilling their role of overseeing surveillance authorities. Despite that there has been significant discussion in the legal community about the constitutionality of NSL provisions in the context of First and Fourth Amendment issues [41, 28, 24, 32, 22, 27], thus far, we are unaware of any comprehensive analysis of NSL reports. Similar to government data, information published by companies shows a similar lack of structure, especially across organizations.

In this paper, we create a consolidated dataset from disparate NSL data sources. Using a combination of manual cleaning and automated scripts, we extract NSL-related information from a broad array of reports from both the public and private sectors. This data enables us to perform an empirical analysis on the use of NSLs over time, as well as to identify

both corroboration and inconsistencies across data sources. We publish the resulting data set to allow other entities to assess the use of NSLs.

Overall, our paper makes three main contributions:

- First, we present the first dataset that consolidates heterogeneous data from government statistics, transparency reports, and published NSLs. With a combination of manual efforts and automated scripts, we produce a normalized dataset that enables comparisons across data sources. We make our scripts and data available to the public to enable future research.
- Second, we observe that the analysis of this data enables insights about the use and reporting of NSLs over time, including the following: NSL requests for non-US persons grew significantly between 2010 and 2015, passing the number of requests for US persons, and remaining popular until today, with significant spikes in frequency in 2015 and 2019. Furthermore, transparency reports show that telecommunication companies (e.g., AT&T, T-Mobile, and Verizon) receive the largest number of NSLs among the reporting companies. Our cross-comparison of transparency reports with the number of NSL requests reported by the government appears to be consistent. Finally, the metadata of published NSLs (after their gag order was lifted) lets us infer their time under disclosure—a metric not reported by the government—which appears to be decreasing since 2010, in response to legal reforms. These analyses suggest that the current reporting mechanisms are effective in preventing the abuse of NSLs.
- Third, we document the challenging task of processing data from various sources due to the absence of standardized and machine-readable formats. This laborious process significantly hinders external and internal auditing efforts. Indeed, in our analysis, we have found various inconsistencies both within and across data sources. Our work demonstrates the need for a unified and more usable dataset for NSLs, which would ultimately lead to better utilization and appreciation of the NSL data available as well as a government with increased transparency.

2 NSL: A Brief Background

Currently, five statutory provisions authorize government agencies to issue NSLs: the Right to Financial Privacy Act (RFPA) [2]; the Electronic Communications Privacy Act (ECPA) [4]; the National Security Act (NSA) [1]; the Fair Credit Reporting Act [7]; and the USA PATRIOT Act [9] (the last of which being the first to codify the term “National Security Letter” explicitly).

In this section, we start by giving a short overview of these five NSL statutes, followed by a summary of efforts to improve transparency around the use of NSLs (notably two amendments to the five NSL statutes that introduced reporting mandates).

2.1 The Five NSL Statutes

The Right to Financial Privacy Act (RFPA) [2] was the first statute used to introduce NSL authorities. As part of a 1986 amendment, the FBI was granted the right to request access to business records from financial institutions (amended in 2003 to include a broader range of organizations). During the same period of time, the second NSL-granting statute, the Electronic Communications Privacy Act (ECPA) [4] was enacted. It provided access (via 18 U.S.C § 2709) to business records (i.e., name, address, length of service, and toll records) of wire or electronic communication service providers for counterintelligence purposes.

In the 1990s, two more statutes were enacted granting further NSL authorities. In 1994, the National Security Act (NSA) was amended to include a procedure for any authorized agency to request a broad array of business records from various organizations for investigating potential document leaks from government employees (codified at 50 U.S.C § 3162 [6]). Later in 1996, the Fair Credit Reporting Act (FCRA) was amended to incorporate the fourth statutory provision for NSLs, which authorized the FBI to access credit agency records in service of national security investigations.

Finally, the PATRIOT Act [9] is the fifth and last NSL-granting statute, which was enacted in 2001 as a response to the terrorist attacks of September 11 [21]. Through amending FCRA, the PATRIOT Act introduced a new procedure for government agencies to access consumer reports from credit report agencies for counterintelligence and counterterrorism investigations. In addition, it made substantial amendments to three of the four existing NSL statutes (RFPA, ECPA, and FCRA), expanding the scope of NSLs and simplifying the administrative approval requirements.

In addition to providing access to sensitive records from an array of private sector institutions, the five NSL statutes initially had no mandatory judicial review, indefinite nondisclosure mandates, and few reporting requirements, making NSLs a powerful tool with limited supervision. Due to concerns from courts and the public, various efforts have been made to increase the transparency of NSL usage and congressional oversight by introducing reporting mandates and weakening nondisclosure requirements, which we discuss below.

2.2 Subsequent Amendments

While the five NSL statutes continue to get amended in subsequent statutes, no new NSL statute was introduced after 2001. We briefly survey important amendments to the five existing NSL statutes after 2001.

2.2.1 RFPA amendment in 2003

In 2003, as part of the Intelligence Authorization Act for Fiscal Year 2004 [8], Congress amended the definition of financial institutions in RFPA to include a much broader range of organizations such as insurance companies and travel agencies.

2.2.2 PATRIOT Act amendments

In 2006, the 109th Congress amended the USA PATRIOT Act with two statutes, in part as a response to earlier judicial reactions to the USA PATRIOT Act [32]. Specifically, two court cases *Doe v. Gonzales* [42] and *Doe v. Ashcroft* [43]¹ raised First and Fourth Amendment issues with NSLs connected to the non-disclosure provisions and the absence of judicial oversight [26]. In the case of *Doe v. Ashcroft*, the district court held that NSLs violated the Fourth Amendment because they authorized “coercive searches effectively immune from any judicial process” [45]. Moreover, the court held that the nondisclosure provisions unconstitutionally restrict free speech as they can prohibit disclosure without providing judicial means to challenge a ban or achieve eventual relief [45, 44].

As a reaction to both cases, and while *Doe v. Ashcroft* was on appeal at the Second Circuit, the 109th Congress amended the USA PATRIOT Act with the USA PATRIOT Improvement and Reauthorization Act [13] and the USA PATRIOT Act Reauthorization Amendments Act [12]. Collectively, besides adding penalties for non-compliance with NSLs or their nondisclosure requirements [11, 10], these two statutes introduced judicial review on the use of NSLs and weakened the nondisclosure requirements. Specifically, the amendments require that the Foreign Intelligence Surveillance Court (FISC) approves NSL requests and adds control mechanisms and a process to ease nondisclosure provisions [11, 10]. Regarding *Doe v. Ashcroft*, the Second Circuit ruled in 2008 that the FBI needs to certify that disclosure of an NSL would lead to statutorily enumerated harms to justify non-disclosure provisions [31]. Last, the amendments call for an audit on use of the NSLs from the Office of the Inspector General (referred to as the OGI reports).

2.2.3 USA FREEDOM Act

In 2015, Congress enacted the USA FREEDOM Act [19], which revised the nondisclosure requirement, introduced an additional reporting requirement, and expressly limited the use of NSL to specifically identified information. Notably, the additional reporting requirement and restrictions on the use of NSLs were direct responses to the bulk metadata collection practice of the NSA [39].

2.3 Reporting and Nondisclosure Mandates of NSLs

Initially, the five NSL statutes had no mandatory judicial review, relaxed reporting requirements, and strict nondisclosure mandates². Given the broad range of data accessible via NSLs, this has unsurprisingly sparked concerns both in the public and, in response to challenges, in the courts [26]. Over time, partially fueled by unfavorable judicial reactions (e.g., *Doe v. Gonzales* [42]) and various public events (e.g., the Snowden disclosures in 2013 [29]),

¹The plaintiff’s name “John Doe” is used in both cases because the gag order prohibited the disclosure of the plaintiff’s identity.

²As an example, the original ECPA and FCRA both prohibited the disclosure to any person.

several efforts have been made to increase the transparency of NSL usage [41, 26]. We highlight the two legislative efforts that are most relevant to this paper.

In 2006, Congress amended the NSL statutes with the USA PATRIOT Improvement and Reauthorization Act [13] and the USA PATRIOT Act Reauthorization Amendments Act [12], establishing a process for judicial review, weakening nondisclosure requirements, and introducing additional reporting requirements. Notably, the USA PATRIOT Improvement and Reauthorization Act added two *unclassified* reporting requirements. Section 119 of the Act initiated an audit by the Office of Inspector General of the Department of Justice to report NSL usage to Congress (“the OIG reports”). However, these reports were each singular occurrences, and the OIG is not expected to publish any further reports unless there is another mandate from Congress. Second, Section 118 required the Attorney General of the Department of Justice to submit annual unclassified statistics about NSLs targeting *US persons* to Congress [14]. These statistics are published as part of the Foreign Intelligence Surveillance Act reports, hence referred to as FISA reports.

In 2015, the USA FREEDOM Act [19] mandated additional reporting, and expressly limited the use of NSLs to explicitly specified information. Notably, it mandates that the Office of the Director of National Intelligence (ODNI) should publish a report that details both the number of NSLs issued, and the number of requests for information contained in those NSLs for the past year (in addition to the FISA reports) [19]. This is published as part of the National Intelligence’s Annual Statistical Transparency Report (ASTR).

Finally, in addition to adding reporting mandates and limiting the scope of NSLs, the USA FREEDOM Act introduced three major changes that relaxed the nondisclosure requirements. First, it allowed companies to report the total number of NSLs received (and customers covered by those NSLs) quantized into bands.³ In practice, a range of companies (particularly in the technology and communications sectors) take advantage of this permission and publish a rough estimate of the number of NSLs they receive as part of their annual transparency reports. We refer to such reports collectively as “transparency reports”. Moreover, the USA FREEDOM Act codified a procedure (18 U.S.C. § 3511) that allows companies to request judicial review of the nondisclosure orders (commonly known as “reciprocal notice”). Lastly, it required the Attorney General to adopt procedures to review nondisclosure orders at appropriate intervals. If disclosure is no longer believed to lead to harms listed in 18 U.S.C. § 2709(c)[5], then the FBI should terminate the nondisclosure order and notify the company. Some companies publish the content of such NSLs (i.e., after the termination of nondisclosure requirements) which can provide useful metadata for analyzing past NSL issuance. We refer to such publications collectively as “company NSLs”.

In summary, three types of NSL reports are currently available to the public:

- OIG, FISA and ASTR reports (collectively referred to as “Government Reports” in Section 3.1) mandated by the USA PATRIOT Improvement and Reauthorization Act

³Public records [26, 17] suggested that this part of the USA Freedom Act was inspired by the pre-existing voluntary agreement reached in 2014 between several technology companies and the DOJ [15].

and the USA FREEDOM Act

- Company transparency reports that include the number of received NSLs
- NSLs disclosed by companies after the NSL’s nondisclosure requirement is lifted

Understanding what the public can learn from these publicly available NSL reports, and the limitations of each data source is the central focus of our work.

3 Data Collection and Trends

Despite various reports published by government agencies (mandated by law) and companies that aim to increase the transparency and auditability of NSL usage, interpreting and auditing these reports can be a tiresome process. Concretely, there is significant manual effort required to identify, collect, clean, normalize and interpret the range of public NSL reports. There exists no single collection of NSL information; finding reports from various entities is not always straightforward and can be time-consuming. Additionally, there is no universal format that all reports follow. They can be formatted differently depending on the agency that published them, or when they were published. Company transparency reports are even less structured, as they are reported voluntarily by different companies. Moreover, this is not merely a matter of format, but the kind of data published may change over time (even within the publication portfolio of a single organization), the semantics of its interpretation (e.g., whether or not one reports that precisely 0 NSLs were received), and the granularity of aggregation. Moreover, data may overlap, such as when NSLs with the same number are served to multiple companies.⁴

One contribution of this paper is the first comprehensive collection of NSL statistics, transparency reports, and published NSL letters scattered across the Internet. In addition, using automated scripts and manual cleaning, we extract all useful information from the NSL reports into a normalized format, enabling several analyses detailed in later sections. We make our data set and scripts public.⁵

Figure 1 provides an overview of the time ranges⁶ and the type of information that we were able to collect from different sources:

- *type 1* (cf. Section 3.1): Multiple government entities publish statistics about NSL usage (*type 1* data): The OIG reports from 2006 [36], 2007 [38], and 2014 [37] contain the sum of all types of ROIs for 2003–2011, and US and non-US ROIs for 2003–2009.

⁴For instance, NSL-19-483160 is served to both Google and Apple.

⁵We publish our data set here: <https://github.com/ucsdysnet/nsl-empirical-analysis>.

⁶These ranges do not necessarily correspond to amendments. For instance, companies only received permission to publish statistics on the number of transparency reports they received after an agreement with the DOJ in 2014 [15]. However, some companies retroactively published their NSL statistics back to 2009 after 2014.

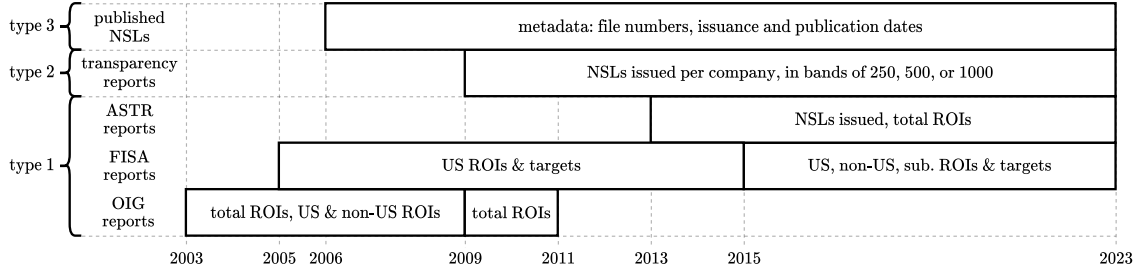


Figure 1: Overview of data available from three types of sources: government statistics (type 1), transparency reports (type 2), and published NSLs (type 3)

FISA reports [35] contain NSL requests and ROIs for US targets since 2005, and data for non-US and subscriber information requests since 2015.

- *type 2* (cf. Section 3.2): Companies may choose to publish transparency reports but are restricted to reporting the number of NSL requests they receive in bands of 250, 500, or 1000. Although non-disclosure orders initially silenced companies, they were allowed to publish such reports in 2014, and some companies retroactively published reports back to 2009.
- *type 3* (cf. Section 3.3): After the gag order was lifted, companies may publish redacted NSL letters that they received. Albeit personally identifiable information is redacted, the letters still contain useful metadata including file numbers and issuance dates. In most cases, companies document when they published an NSL.

Broadly speaking, the reports we collect come from three sources. First, government reports such as statistics published by the FBI. Second, transparency reports from companies that contain the number of NSL requests they received. Third, redacted NSLs published by companies after their gag order was lifted. Below, we detail how we collect and clean data from each source and observe general trends.

3.1 Government Reports

NSL statistics collected from various government reports [34, 35, 36, 38, 37] include the number of NSL requests for US and non-US persons, as well as subscriber information for any person, including both US and non-US persons. Figure 2 depicts this data, enumerating the number of targeted individuals for every type of NSL request. A single NSL can contain multiple requests for information, or ROIs, as long as they are in the context of a single investigation. As well, we note that the FBI may serve multiple NSLs for the same person under different statutes and to different entities (e.g. to collect information on telephone

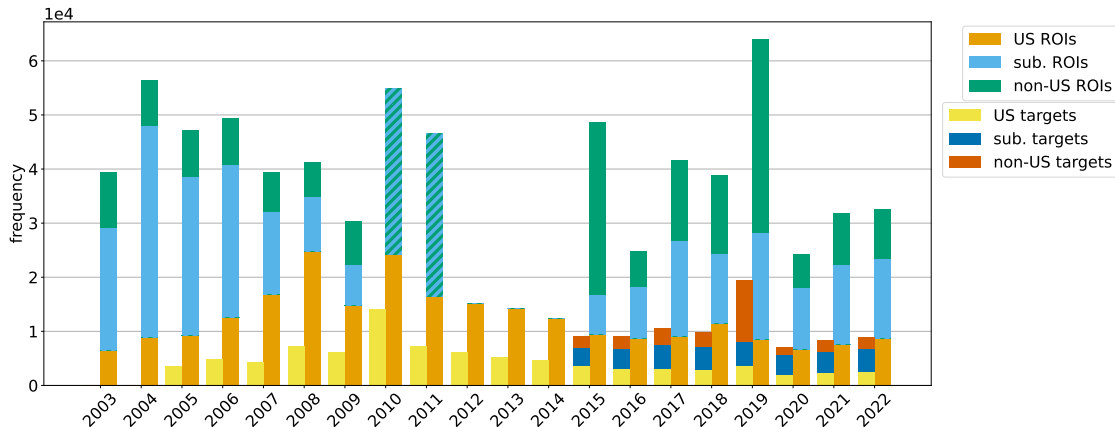


Figure 2: Number of requests and targets for different NSL types from 2003 to 2022. For every year, the left bar shows the type of targets in NSL requests, starting from 2005: US persons, non-US persons, or letters for subscriber information (for which the target’s nationality is unknown a priori). The right bar is aggregated from various sources and shows the types of Requests of Information (ROI; an NSL can contain multiple of them), split into the same categories.

numbers, email addresses, and financial records) and thus the number of NSL requests can be greater than the number of targets.

The FISA reports [35, 3] started to mandate reporting NSL requests for US persons in 2003. However, the format of the reported data changed multiple times since then due to new regulations and policies. In 2005, section 128 of the USA PATRIOT Improvement and Reauthorization Act added the requirement to report the number of ROIs for US persons made with NSLs [14]. Starting in 2015, the USA FREEDOM Act [19] additionally required the reporting of NSL requests for non-US persons and NSL requests for subscriber information. The ASTR [34] published by the ODNI reports the total number of issued NSLs and ROIs.

For 2003–2009, we found reports of the OIG [38, 37] to contain statistics (that were redacted in an early version but disclosed in the 2014 revision [37]) from reports of the FBI to Congress under FISA on the number of NSL requests for non-US persons (shown in yellow in Figure 2). The OIG reports also include the total number of NSL requests until 2011, from which we can infer the number of non-US ROIs for 2003–2009 and the combined number of non-US ROIs and subscriber information requests for 2010 and 2011. Unfortunately, the OIG reports do not contain any additional information about non-US NSL targets. Thus, for the period from 2012 to 2014, we only know statistics about NSL requests and targets pertaining to US persons due to the lack of declassified information.

Furthermore, Figure 2 indicates that the number of NSL requests for US persons grew

steadily until 2010, and then eventually decreased back to the level of 2003. Partial data from OIG reports [38, 37] suggests that NSL requests for non-US persons were in the minority before 2010, but became more common between 2010 and 2015 and still remain popular. Additional data from the FISA reports shows that after 2015, the US-targeted NSL requests represent a smaller portion of all requests, although a significant portion of requests are made for subscriber information of any person.⁷

3.2 Transparency Reports

Transparency reports provide the perspectives of individual companies on NSL issuance. In the United States, there is neither a legal requirement nor a standardized format for publishing transparency reports. Moreover, the government restricts the reporting of the number of NSL requests to bands of 250, 500, or 1000 as defined by 50 U.S.C. § 1874 [18]. These restrictions and lack of structure cause the number, frequency, content, and accessibility of transparency reports to vary wildly by company.

We collect transparency reports published by 55 US-based private companies (listed in Appendix A.1), of which 41 provided data related to NSLs. While some of the companies have all of their transparency reports easily archived and downloadable as CSV or PDF files, other reports are hidden in blog posts or support forum answers, posted as low-quality images, or have broken links. In addition, we had to exclude the reports of seventeen companies as they only released aggregated numbers for all national security requests they received, mixing NSLs with other FISA requests (e.g., for electronic surveillance). Clearly, our dataset is not comprehensive and only provides lower bounds on the NSL usage. It is possible that we missed some transparency reports in our search, and the data of companies that choose not to disclose this information is absent.

Despite the incomplete data set and that companies report the number of NSLs in coarse bands, we can still compare the sum of lower respective upper bounds with the total number of issued NSLs from ASTR as shown in Figure 3. We find that the number of NSL requests from ASTR exceeds the cumulative upper bound from transparency reports, indicating that we are indeed missing a significant fraction of company reports for NSLs. Although the company transparency records are not a comprehensive data source, we observe that after 2013, the mean of the bands strongly correlates with the contemporaneous number of issued NSLs (with a Pearson coefficient of 0.77). This indicates that the two data sources are, at the very least, not inconsistent. However, more fine-grained data or more standardized transparency reports are needed to support stronger conclusions.

The growing distance between cumulative lower and upper bands in Figure 3 from

⁷It is tempting to relate fluctuations in NSL requests to global political events. For instance, the peak in 2019 for non-US ROIs may be due to increased investigations into digital attacks to perform economic espionage, which dominated the FBI’s counterintelligence program in 2019 according to FBI director Christopher Wray [46]. However, it is impossible to corroborate such a hypothesis using only the data available here.




Figure 3: The highlighted area shows the accumulated transparency report ranges, computed as the sum of lower and upper limits reported by some companies. We compare them to the line of government-reported number of issued NSLs from the ASTR.

2009 to 2013 is due to the growing number of published transparency reports. Initially, only Google and Twitter published reports. There was a sharp spike in the number of companies publishing reports around 2013 and 2014, presumably in response to the Snowden revelations drawing increased attention to transparency. Some companies stopped issuing reports since then, but the decrease in reports in 2022, which is reflected in the decreased width of the band, is likely an artifact of delayed report publishing by companies.

Diving into the reported ranges of NSLs themselves, we find that 13 out of 41 reporting companies explicitly stated that they had never received any NSLs. The vast majority of companies report the lowest band for NSL requests (i.e. 0-249, 0-499, or 0-999). Telecommunication companies, such as AT&T, T-Mobile, and Verizon, received the most NSLs, with up to 1000-2000 requests each year. Apple and Google, both large producers of phones and mobile operating systems, were the only other companies reporting a higher range of NSL requests. We surmise that telecommunications and related companies may lead these statistics because they collect valuable metadata for investigations.⁸

⁸In comparison, companies like Adobe are less likely to be utilized by malicious actors in a way that produces useful metadata for investigations.

3.3 Company NSLs

The third and last data source are NSLs themselves, voluntarily published by their recipients⁹ after the nondisclosure requirement (colloquially “gag order”) has been lifted. These letters request specific account information (called ROIs) and have a gag order attached, restricting the receiving party from publishing the orders or even discussing their existence. The most sensitive parts of the letters are redacted, including user identities and confidential information such as Social Security Numbers. However, unredacted information includes the types of information requested (e.g. email metadata or credit records) and types of user identifiers (e.g. email addresses, account numbers, addresses, or names). Additionally, we found file numbers and publication and issuance dates (which are often not redacted) to be useful metadata, as well as that the number of redacted lines indicates the ROI volume.

Google and Apple are the largest publishers of NSL letters in our data set with 272 and 45 letters, respectively. Appendix A.2 lists the full composition of our data set. We only see a fraction of all existing letters. Between Jan 1, 2015 and Dec 31, 2020¹⁰, companies published only 0.3% of all issued NSLs. We discuss interesting trends that can still be observed from the available data in the remainder of this section.

We use the number of days between the issuance and publication dates as an estimator for the duration of the gag order. The company-internal administrative delay between lifting a gag order and publishing the letter may add some noise to this metric.

The scatter plot in Figure 4 shows for all public NSLs, how many years after their issuance that companies publish them. We observe that, except for two NSLs, all letters were published after June 2015, when the USA FREEDOM Act [19] amended NSL reporting and nondisclosure regulations, limiting gag orders to the duration for which publication would interfere with national security interests or ongoing investigations, and simplifying the disclosure process. This supports, that in the absence of public scrutiny, the gag orders before 2015 were easier to enforce and very rarely lifted. After Congress relaxed the nondisclosure requirements, publishing NSLs became possible for companies. It appears that few companies retroactively published NSLs that they received before these legal reforms.

However, two important caveats need to be mentioned in connection with Figure 4. First, we only see 0.3% of all NSLs, and it is unknown whether they are a representative sample for the distribution of the remaining 99.7% of NSLs. Second, data since 2013 is increasingly skewed towards shorter gag order times, as any letter that is in the “future dates” area will only be published in the future. For instance, letters issued in 2022 with a gag order time of four years will only be published in 2026.

An additional challenge in interpreting NSL data is the unknown bias stemming from possibly selective publishing, due to which gag order the government decides to lift and which letters companies voluntarily publish. Companies have no obligation to publish

⁹Two avenues for publishing NSLs are the reciprocal notice procedure and the termination procedure.

¹⁰We pick these years such that they are after the gag order was weakened but not too recent to avoid too much skew because of non-disclosure orders that are still preventing companies from publishing the letters.




Figure 4: This graph shows how many years it took to publish NSLs after they were issued. Most letters were only published after the USA FREEDOM Act added procedures to challenge gag orders. The gag order time appears to decrease. However, the distribution of 99.7% of the NSLs is unknown. Moreover, the gag order time is biased towards shorter times, as letters from recent years with longer times have not been published yet.

NSLs, and we have little insight into their internal reasoning for whether or not to disclose NSLs: there might be a selection bias in the volume and type of NSLs they publish.

4 Data Inconsistencies

This section discusses inconsistencies in the reported NSL data when comparing different sources. Furthermore, we point out open questions that oversight bodies may consider investigating.

4.1 Diverging Requests for Information Counts

The Office of the Director of National Intelligence (ODNI) started to publish an Annual Statistical Transparency Report (ASTR) in 2014 [34] including the number of issued NSLs and—similar to FISA [35]—NSL requests. The ASTR reports were a reaction to the first leaked documents by Snowden [40], ordered in June 2013 [16] to increase transparency.

Year	US ROIs	Non-US ROIs	Subscriber ROIs	FISA ROIs	ASTR ROIs
2015	9418	31863	7361	48642	48642
2016	8727	6651	9432	24810	24801
2017	9006	14861	17712	41579	41579
2018	11454	14481	12937	38872	38872
2019	8557	35848	19601	64006	63466
2020	6670	6187	11368	24225	24225
2021	7607	9486	14732	31825	39214
2022	8587	9103	14927	32617	32617

Table 1: The number of reported Request of Information (ROI) according to different sources. We sum the first three columns to derive the total number of ROIs reported under FISA and compare them with the ASTR counts reported by the ODNI.

The columns “FISA ROIs” and “ASTR ROIs” of Table 1 show the reported number of NSL requests. The “FISA ROIs” column is the sum of the reported NSL requests for information for US persons, non-US persons, and subscriber information under FISA in the previous three columns. Before 2015, the FISA numbers did not include non-US persons and subscriber information and can, therefore, not be compared to the ASTR numbers. However, we were unable to find any public statement (or observation) of the differing numbers in 2016, 2019, and 2021. It seems unlikely that there was a difference in the definition or counting of ROIs as they match exactly for the other years. While the numbers for 2016 and 2019 differ by less than 1%, there is an almost 19% deviation for 2021.

4.2 NSL Metadata

File numbers are part of the scarce metadata of NSL letters. Figure 5 plots the file numbers (y-axis) over the issuance date of the letter, for all NSLs that were published by companies after their gag order was lifted. It also plots the cumulative number of total NSL requests reported by ASTR, and the number of issued NSLs¹¹. The dashed red lines are the linear regression of the respective data. NSLs were issued since their authorization in 1978 and, therefore, the reporting that started in 2005 respectively 2013 has an unknown offset defined by the sum of previous NSL counts of the same type.¹² For ease of comparison with our hypothesis, we set the y-axis offset for these lines to the last previously known file number (for instance, for data starting in 2013, the last file number of an NSL from 2012 defines the y-axis offset).

We note that the file numbers are increasing roughly monotonic where 22% of the numbers deviate and show slight decreases of 0.4% on average. It seems to be a reasonable

¹¹The x-axis labels mark the start of a year and we place the cumulative count at the end of a year. For example, the number of NSLs issued in 2014 is added to the cumulative count on Dec 31, 2014. Hence, it appears closer to the label of 2015.

¹²Early issues with the FBI tracking system add further uncertainty. The OIG reports that NSLs have been issued under the case file numbers of another division [37].

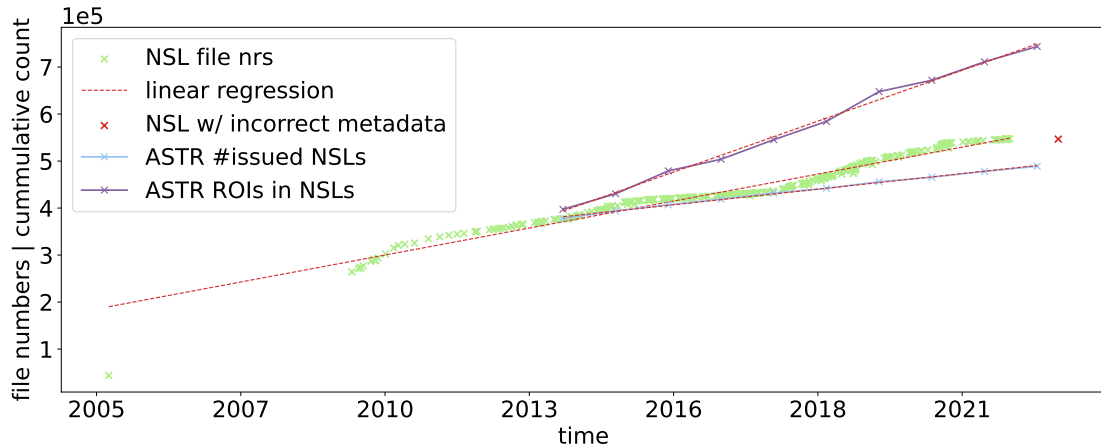


Figure 5: The file numbers of NSLs plotted over time show a monotonically increasing sequence of points. We compare them to the cumulative ROI and NSL counts published in the ASTR, after adjusting their y-axis offset to match the number of files in 2014. File numbers with red crosses have erroneous issue dates reported by companies.

hypothesis that NSL letters get assigned consecutive numbers by the FBI before their issuance (with some being issued faster than others afterward). The blue line of issued NSLs as reported by ASTR grows slower than the number of files. It is conceivable that some of the file numbers are assigned to NSLs that the FBI prepared to issue but withdrew before serving them to a company.¹³

Qualitatively, we observe that there appears to be some correlation between the cumulative number of ROIs and the number of files, delayed by a few months. The file numbers have two steeper increases, one from 2014–2015 and another from 2018–2019. The ROI counts show similar characteristics in 2015 and 2019 (recall, the cumulative counts are reported closer to the ticks in 2016 and 2020). These increases in the ROI counts are mainly caused by the spikes in the NSL requests for non-US persons, as Table 1 shows. However, we observe the possible anomaly that these steeper increases in file numbers are not reflected in the reported number of NSLs, despite that they appear to be present in the cumulative ROI counts. The expectation would have been that the file numbers follow trends in the number of issued NSLs and not the ROIs since a single NSL can contain multiple ROIs but should have only one file number. One possible explanation could be a difference in the counting methodology of ROIs for non-US persons. However, without internal insight into the as-

¹³A reason for more file numbers than issued letters would be if an investigation takes an unexpected turn that makes the NSL obsolete. For instance, when the FBI extracts the requested data directly from the laptop of the suspect with forensic methods, then there is no reason to serve the NSL and it might not be counted in the annual reports despite that it was already assigned a file number in the issuing process.

signment of file numbers, we are limited to observing these unexplainable inconsistencies, which negatively affects the confidence in the accuracy of the reported data.

5 Acknowledgments

The authors would like to thank the anonymous reviewers of CSLAW'24 for their insightful comments and suggestions. Furthermore, we thank Alisha Ukani and Stewart Grant for supporting us in interpreting our data with their data processing and visualization expertise.

6 Conclusion

Over the years, multiple efforts have been made to increase the transparency and congressional oversight on the use of National Security Letters (NSLs). Congress has amended the five NSL statutes several times, introducing mandatory reporting requirements, dozens of companies now publish NSL data (both individual requests and aggregated counts) while pushing for permission to publish yet finer-grained information [33]. However, despite these efforts, it is still challenging to assess how much transparency has really been provided as the data collection is challenging. No single data source provides a comprehensive view of how NSLs are used, and the manual effort entailed in collecting, extracting, normalizing, and analyzing disparate reports is an implicit obstacle to a more holistic analysis. This, in turn, results in uncertainty about what the public can and cannot understand about the use of NSLs, or the effectiveness of the amendments to NSL statutes.

We address these issues by gathering and publishing the first comprehensive collection of publicly reported data on NSLs, leveraging three sources: government statistics (such as reports by the FBI to Congress under FISA [35]), company transparency reports, and NSLs themselves, that were published after their nondisclosure order was lifted.

Our curated data set allows us to empirically analyze the use of NSL letters and compare trends across different sources. For example, from this data, we show a rough (albeit incomplete) agreement between corporate transparency reports and government data. Further, we use the small fraction of published NSL letters to infer the gag order time, which peaked in 2012 and has decreased since. Finally, we highlight that while the number of Requests for Information (ROIs) for US persons increased rapidly until 2010, it has since been superseded by the large but varying number of ROIs for non-US persons.

We advocate for publishing NSL statistics in a standardized and easily processable format (e.g., following the suggestions from [23]). This would simplify analyzing the NSL usage and trends that we showed in this paper to be insightful, and support Congress and the public in fulfilling their role to audit the government agencies. Furthermore, it would allow for earlier detection of internal and external inconsistencies, such as the discrepancy in FISA and ASTR reports, as well as the side-channel information from the file numbers that is inconsistent with the reported number of issued NSLs.

References

- [1] National Security Act of 1947. <https://govtrackus.s3.amazonaws.com/legislink/pdf/stat/61/STATUTE-61-Pg495.pdf>, July 1947. visited on June 12, 2023.
- [2] 12 U.S. Code Chapter 35 – Right to Financial Privacy Act (RFPA). <https://www.law.cornell.edu/uscode/text/12/chapter-35>, November 1978. visited on June 9, 2023.
- [3] 50 U.S. Code § 1807 – report of electronic surveillance. <https://www.law.cornell.edu/uscode/text/50/1807>, October 1978. visited on June 9, 2023.
- [4] 18 U.S. Code § 2701 – unlawful access to stored communications. <https://www.law.cornell.edu/uscode/text/18/2701>, October 1986. visited on June 12, 2023.
- [5] 18 U.S. Code § 2709 – counterintelligence access to telephone toll and transactional records. <https://www.law.cornell.edu/uscode/text/18/2709>, October 1986. visited on June 12, 2023.
- [6] 50 U.S. Code Subchapter VI – Access to Classified Information. <https://www.law.cornell.edu/uscode/text/50/chapter-44/subchapter-VI>, October 1994. visited on June 9, 2023.
- [7] 15 U.S. Code § 1681u – Disclosures to FBI for counterintelligence purposes. <https://www.law.cornell.edu/uscode/text/15/1681u>, January 1996. visited on June 9, 2023.
- [8] H.R. 2417 (108th): Intelligence Authorization Act for Fiscal Year 2004. section 374. <https://www.govtrack.us/congress/bills/108/hr2417/text>, October 2001.
- [9] Uniting and Strengthening America by Providing Appropriate Tools Required to Intercept and Obstruct Terrorism (USA PATRIOT ACT). <https://www.congress.gov/107/plaws/publ56/PLAW-107publ56.pdf>, October 2001.
- [10] 18 U.S. Code § 1510 – obstruction of criminal investigations. <https://www.law.cornell.edu/uscode/text/18/1510>, March 2006.
- [11] 18 U.S. Code § 3511 – judicial review of requests for information. <https://www.law.cornell.edu/uscode/text/18/3511>, March 2006.
- [12] S. 2271 (109th): USA PATRIOT Act Additional Reauthorizing Amendments Act of 2006. <https://www.govtrack.us/congress/bills/109/s2271/text>, March 2006.
- [13] USA PATRIOT improvement and reauthorization act of 2005. <https://www.congress.gov/109/plaws/publ177/PLAW-109publ177.pdf>, March 2006.
- [14] USA PATRIOT improvement and reauthorization act of 2005. section 128. <https://www.congress.gov/109/plaws/publ177/PLAW-109publ177.pdf>, March 2006.

- [15] Letter from James M. Cole. <https://www.justice.gov/iso/opa/resources/366201412716018407143.pdf>, January 2014. visited on September 29, 2023.
- [16] Statistical Transparency Report Regarding use of National Security Authorities. https://www.intelligence.gov/assets/documents/702%20Documents/statistical-transparency-report/National_Security_Authorities_Transparency_Report_CY2013.pdf, June 2014.
- [17] H.Rept 114-109 (2015). <https://www.govinfo.gov/content/pkg/CRPT-114hrpt109/pdf/CRPT-114hrpt109-pt1.pdf>. visited on June 12, 2023.
- [18] 50 U.S. Code § 1874 – public reporting by persons subject to orders. <https://www.law.cornell.edu/uscode/text/50/1874>, June 2015. visited on June 15, 2023.
- [19] Uniting and Strengthening America by Fulfilling Rights and Ensuring Effective Discipline over Monitoring (FREEDOM ACT) act of 2015. <https://www.congress.gov/bill/114th-congress/house-bill/2048/text>, June 2015. visited on June 9, 2023.
- [20] The opacity of National Security Letters. <https://www.cornellpolicyreview.com/the-opacity-of-national-security-letters/?pdf=2352>, April 2016. visited on September 29, 2023.
- [21] Patriot act. https://en.wikipedia.org/wiki/Patriot_Act. visited on June 12, 2023.
- [22] Hannah Bloch-Wehba. Process without procedure: National Security Letters and First Amendment rights. *Suffolk UL Rev.*, 49:367, 2016.
- [23] Joe Calandrino. Government datasets that facilitate innovation. <https://freedom-to-tinker.com/2010/03/01/government-datasets-facilitate-innovation/>, March 2010.
- [24] Rachel Dallal. Speak No Evil: National Security Letters, gag orders, and the First Amendment. *Berkeley Tech. LJ*, 33:1115, 2018.
- [25] data.gov. Data.gov. <https://data.gov/>. visited on Feb 26, 2024,.
- [26] Charles Doyle. National Security Letters in foreign intelligence investigations: Legal background, July 2015. RL33320 Version 18.
- [27] Patrick P Garlinger. Privacy, free speech, and the patriot act: First and Fourth Amendment limits on National Security Letters. *NYUL Rev.*, 84:1105, 2009.
- [28] Ursula Gorham-Oscilowski and Paul T Jaeger. National Security Letters, the USA PATRIOT Act, and the Constitution: The tensions between national security and civil rights. *Government Information Quarterly*, 25(4):625–644, 2008.

- [29] Glenn Greenwald. NSA collecting phone records of millions of Verizon customers daily. <https://www.theguardian.com/world/2013/jun/06/nsa-phone-records-verizon-court-order>, June 2013.
- [30] GSA.gov. Open Government Initiative. <https://www.gsa.gov/governmentwide-initiatives/open-government-initiative>. visited on Feb 26, 2024,.
- [31] Jon O. Newman. 549 F.3d 861 (Second Circuit 2008). Doe v. Mukasey (formerly known as Doe v. Ashcroft). https://www.aclu.org/sites/default/files/pdfs/safefree/doevmukasey_decision.pdf, August 2008.
- [32] Andrew E Nieland. National Security Letters and the amended PATRIOT Act. *Cornell L. Rev.*, 92:1201, 2006.
- [33] United States Court of Appeals for the Ninth Circuit. Twitter, Inc. V. Merrick Garland, et al, no. 20-16174 (9th cir. 2023). <https://storage.courtlistener.com/recap/gov.uscourts.ca9.323074/gov.uscourts.ca9.323074.54.1.pdf>, March 2023.
- [34] Office of Civil Liberties, Privacy, and Transparency. Annual Statistical Transparency Report. <https://www.intel.gov/annual-statistical-transparency-report>. visited on June 9, 2023.
- [35] U.S. Department of Justice. FISA Reports. <https://www.justice.gov/nsd/fisa-reports>. visited on June 9, 2023.
- [36] Office of the Inspector General. A review of the Federal Bureau of Investigation’s use of National Security Letters in 2006. <https://oig.justice.gov/sites/default/files/legacy/special/s0803b/final.pdf>, March 2008.
- [37] Office of the Inspector General. A review of the Federal Bureau of Investigation’s use of National Security Letters. <https://oig.justice.gov/reports/2014/s1408.pdf>, August 2014.
- [38] Office of the Inspector General. A review of the Federal Bureau of Investigation’s use of National Security Letters, March 2007 (january 8, 2016 version). <https://www.oversight.gov/sites/default/files/oig-reports/o1601b.pdf>, January 2016.
- [39] Office of the Inspector General. U.S. Department of Justice. A Review of the FBI’s Use of Section 215 Orders for Business Records in 2012 through 2014. <https://oig.justice.gov/reports/2016/o1604.pdf>, September 2016. visited on June 10, 2023.
- [40] Betsy Reed. NSA collecting phone records of millions of Verizon customers daily. <https://www.theguardian.com/world/2013/jun/06/nsa-phone-records-verizon-court-order>, June 2013.

- [41] Kevin J Schrop. Your cooperation is greatly appreciated: The Fourth Amendment, National Security Letters, and public-private data sharing. *Penn St. L. Rev.*, 122:849, 2017.
- [42] D. Connecticut United States District Court. 386 F.Supp.2d 66 (2005). Doe v. Gonzales. <https://www.leagle.com/decision/2005452386fsupp2d661441>, September 2005.
- [43] S.D. New York United States District Court. 334 F.Supp.2d 471 (2004). Doe v. Ashcroft. <https://www.leagle.com/decision/2004805334fsupp2d4711766>, September 2004.
- [44] S.D. New York United States District Court. 334 F.Supp.2d at 476. Doe v. Ashcroft. <https://www.leagle.com/decision/2004805334fsupp2d4711766>, September 2004.
- [45] S.D. New York United States District Court. 334 F.Supp.2d at 506. Doe v. Ashcroft. <https://www.leagle.com/decision/2004805334fsupp2d4711766>, September 2004.
- [46] Christopher Wray. The FBI and the national security threat landscape: The next paradigm shift. <https://www.fbi.gov/news/speeches/the-fbi-and-the-national-security-threat-landscape-the-next-paradigm-shift>, April 2019.

A Data Sources

In this section, we list the companies for which we can find transparency reports as well as companies that publish individual NSLs.

A.1 List of Transparency Report Companies

We used transparency reports from the following companies to collect our data:

- 23andMe
- Adobe
- Airbnb
- Amazon
- Apple
- AT&T
- Cisco
- Cloudflare
- Coinbase
- Comcast
- cPanel
- Credo
- Discord
- DreamHost
- Dropbox
- eBay
- Etsy
- Evernote
- Facebook
- GitHub
- Google
- IBM
- Kickstarter
- Lantern
- Let’s Encrypt
- LinkedIn
- Lookout
- Lyft
- Mapbox
- Medium
- Microsoft
- nest
- Netflix
- Pinterest
- Reddit
- Ring
- Slack
- Snapchat
- Sonic
- Sonos
- SpiderOak
- T-Mobile US
- TikTok
- Tumblr
- Twilio
- Twitch
- Twitter
- Uber
- Verizon
- Virtru
- Wickr
- Wikimedia
- Word Press
- Yahoo
- Zoom

A.2 Number of NSLs per Company

Table 2 shows the number of National Security Letters (published after their gag order was lifted) contained in our data set.

Company	# of published NSLs
Google	272
Apple	45
Facebook (now Meta)	15
Twitter	11
Automattic/WordPress	5
Yahoo	3
CREDO Mobile	2
Internet Archive	2 (redacted file numbers)
Twilio	2
Cloudflare	1
Library Connection	1
Microsoft	1

Table 2: Statistics on NSLs published by companies