DOI:10.1145/2663341

Despite continuing media coverage, the public's privacy behaviors have hardly changed.

BY SÖREN PREIBUSCH

Privacy Behaviors After Snowden

MASS MEDIA HAVE been reporting on global-scale state surveillance following former NSA contractor Edward J. Snowden's exposure of PRISM in June 2013. Extensive, continuing news coverage makes this revelation a natural experiment. In a longitudinal study from May 2013 to January 2014, I examined the immediate and longer-term effects on Web use in the U.S. I combined evidence of privacy self-protection and behaviors indicating an interest in privacy. Users' interest rose after the PRISM revelation but returned to and even fell below original levels despite continuing media coverage. I found no sustained

growth in the user base of privacy-enhancing technologies (such as anonymizing proxies). The worldwide public revelation of PRISM affected individuals' interest less than other breaking news concerning sports and gossip. My results challenge the assumption that Web users would start to care more about their privacy following a major privacy incident. The continued reporting on state surveillance by the media contrasts with the public's quickly faded interest.

Many details about PRISM and related cyberintelligence initiatives, including Tempora and XKeyScore, are still unknown by the public,⁸ though news coverage following June 6, 2013 was thorough and steady.⁵ For the first time since 2009, the U.S. national newspaper *USA Today* reported on government surveillance on its front page,⁷ and *The Washington Post* ran a front-page article on the invasion of privacy for the first time in 2013,¹⁶ prompting contentious discussion of the revelations among privacy experts, as well as politicians and advocates, worldwide.

Mass media is able to set the public agenda through increased coverage of a topic.¹² However, little is known about how media salience of privacy might steer public opinion and behavior toward caring more about privacy. Previous media analyses suffered from unreliable and temporarily sparse measures of privacy concern¹⁷ or simply ignored the public's response.²¹ The PRISM

» key insights

- In June 2013, Snowden revealed worldwide government surveillance, a privacy shock also representing a natural experiment in the study of privacy behaviors.
- The effect on the public's interest in privacy and behaviors was limited and short-lived despite long-running, in-depth media coverage.
- High-resolution data from multiple sources shows the greatest privacy meltdown of our time had only a small impact on Web users beyond debates among journalists and academic researchers.

revelations blurred the distinction between privacy in state/citizen and company/customer relationships, as corporations were revealed to be accomplices in government surveillance. Some citizens reacted by modulating their consumption of Web-based services.

Mistrust by German Internet users of government and corporate data processing increased by nine percentage points in 2013¹ over 2011, but only a minority reported they had changed how they manage their personal data.² Google Trends suggested a chilling effect that varied by country on users' propensity to issue search queries that might "get them into trouble" with the U.S. government.¹¹ However, such public-opinion snapshots derived through ad hoc surveys were unable to capture consumers' volatile and ephemeral reactions. Annual polls provide only coarse temporal resolution, when continuous records of behaviors are needed to measure the effect of an event like the PRISM revelation. Privacy surveys also suffer from such shortcomings as respondents' lack of commitment, leading to post hoc rationalizations and socially desirable but unreliable answers.

Acquiring data on user behavior is difficult,⁴ and companies rarely disclose usage statistics of their Web-based services. When the Web search engine DuckDuckGo, which advertises its superior privacy practices, attributed a rise in its daily queries to the PRISM revelation, it did not include user counts.^{3,20}

PRISM also represents an opportunity to carry out a Web-scale longitudinal study of privacy behaviors, and this article offers the first high-resolution analysis using primary sources of data. I combined indicators from multiple primary data sources to explore the evolution of consumers' privacy behaviors. Privacy self-protection can be observed directly, and information-seeking activities can be interpreted as indicators of an interest in privacy.

The mainstream media reported the first details about the NSA's state surveillance programs on June 6, 2013, hereafter called "PRISM day"; the three preceding weeks are the reference period. Over the following six and 30 weeks, respectively, I analyzed the immediate extended reactions by daily

The continued reporting on state surveillance by the media contrasts with the public's quickly faded interest.

measurements of several key metrics:

Information seeking. Web search and browsing, as measured through visits to the privacy policy posted at Microsoft.com, one of the companies mentioned in the early media coverage, and also through visits to privacy-related Wikipedia pages, and privacy-related webpages in general;

Privacy-enhancing configurations. Privacy settings made in the Internet Explorer Web browser, as well as adoption of Tor and the Anonymox Firefox add-on, both connection-concealing mechanisms; and

New webpages. As a control variable, media coverage in terms of number of new webpages created around the topic of government surveillance and declines (following the initial reports) on the existence of PRISM.

Methodology

I analyzed user behavior for U.S. English-language consumers. The U.S. is the home of PRISM and related programs, as well as where the first related media coverage appeared, an obvious yet deliberate limitation of my study. Future work will examine differences across countries, but for this article, I avoid confounding factors of language and culture. The U.S. is the one country for which the most comprehensive data is available for Web searches, visits to corporate privacy policies and Wikipedia pages, and use of Tor/Anonymox. I considered the totality of Web users within the U.S., not a sample. As an exception, I examined general page-visit data and Internet Explorer privacy settings on a global sample of users, owing to a lack of country information. I thus made any comparisons with care.

I recorded Web search behavior for manually chosen PRISM-related keywords (see Figure 1) I then cross-checked against privacy-related keywords inferred from query reformulation, a standard technique in studying Web search behavior. I favored semi-manual selection over a necessarily retrospective data-driven approach that would deny the post-Snowden sea change on the public perception of privacy. U.S. user counts were from the totality of spell-corrected queries on Microsoft's search engine Bing (such as "privcay," which also counted).

Fluctuations in query-term popularity happen naturally in any year; for instance, daily search volume for the term “Wikipedia” varied by up to 33% during the reference period, but its relative share varied by only 0.01 percentage point. To accommodate for varying total search volumes, namely weekday/weekend, differences in general popularity, and varying search patterns across users, I recorded changes in the proportion of users initiating a given query relative to all users, indexed to the reference period. I used weeklong binning to even out weekday/weekend effects, so start days did not have to be aligned when comparing responses to different events.

I measured webpage visits to Microsoft’s privacy statement at <http://www.microsoft.com/privacystatement/> by its four main sections—Bing, Microsoft.com, Xbox, and other Microsoft products—that together accounted for more than 99% of all page views. I also measured general browsing to webpages on PRISM-related topics. I counted page visits if the URL or the title of the page included the keyword. I collected data from a sample of users consenting to share their browsing history, possibly leading to bias toward less-privacy-concerned users.

Tor users refresh their list of running relays on a regular basis. These requests to directory mirrors are useful for estimating the number of Tor users.²³ In line with the other analyses, I focused on U.S. users, who contributed the largest national user share (18%) in the reference period. Anonymox users install a Firefox add-on for easy anonymous browsing.¹⁵ Anonymox is among the top-10 “privacy and security” add-ons in the U.S. and the only one to advertise itself as avoiding government threats to privacy. I counted active users, again for the U.S., by daily update pings.

I assessed media coverage by counting online documents mentioning three keywords: Snowden, PRISM, and surveillance. Discovery was through a subset of the Web index maintained by Bing,¹³ with counts relative to all new documents. The term “privacy” produced false positives and was excluded, as many websites must now include (by law) a link to their “privacy” policy or statement.

Web Search Behavior Post-PRISM

Web search queries represent a lens into user behavior,⁴ serving as a proxy for user interests. For example, the Flu Trends analysis uses Web search as a precursor of medical conditions and behavior.⁶ The inaccuracies subsequently found by others in the Flu Trends models are due mainly to positive feedback loops⁹ that do not apply to my analysis. I measured interest in PRISM-related topics by the number of queries issued about them.

The composite measure of search behavior based on all relevant, automatically discovered “privacy”-related keywords showed neither an immediate ($p = 0.14$) nor longer-term trend ($p = 0.84$, F-test over linear regression).

Short-term evolution. Considering individual topics, Web search queries about the NSA, surveillance, and the government saw a modest spike on PRISM day, June 6, 2013 (+0.08% to +0.02%, all search volume numbers in percentage points). Searches for “Snowden” increased when that name was revealed two days later, June 8. Snowden and NSA were the only terms that continued to attract elevated search volume over the six weeks following PRISM day, with a steady downturn for “NSA.” There was only a slightly increased interest

in “privacy,” with a maximum increase of 0.003% during the week following PRISM day. Searches for PRISM itself rose by 0.01% in the same week but fell to the original level afterward, as in Figure 1. The CIA, another U.S. federal agency involved in collecting foreign intelligence, did not attract any more queries, and related search volume decreased after PRISM day.

Whereas the celebrity Snowden remained popular among search users ($p < 0.0001$, t-test, before versus after PRISM day), privacy could not maintain a larger user base ($p = 0.4$). This observation was corroborated when I examined search for news results only. The average user base searching for “Snowden” rose by 0.007% when comparing before and after PRISM day, whereas news searches for privacy were unaffected, with less than 0.00001% change.

Searches for privacy-enhancing technologies saw a very small increase, with general encryption (at most +0.001% in week one), Tor, a system for users to conceal with whom they communicate (+0.002%, same week), and PGP, a system for encrypted messaging (+0.0003%, same week).

Longer-term evolution. Interpreting the longer-term evolution of search be-

Figure 1. Changes in search-term volume after PRISM day, June 6, 2013, in percentage points relative to reference period.

havior is more difficult, as the increasing time from the reference period introduced seasonal variations and other influential events. By week 11 following PRISM day, all terms I considered had a search volume elevated by no more than 0.02%, as in Figure 1. Snowden and NSA continued to spike in September/October and November/December 2013, respectively, up to +0.06%. I ob-

served no notable increases for PRISM, surveillance, or CIA.

Long after PRISM day, “privacy” spiked in weeks 19 and 28, with 0.03 and 0.05 percentage points, respectively, above the reference period, reaching levels unattained directly following the revelations. Besides continued media coverage about government surveillance, those weeks coincided

with media reports about Facebook removing a privacy setting and Google removing privacy enhancements in Gmail and Android. Likewise, the spike in “government” searches from week 15 to 21 (up to +0.50%) can be attributed to the U.S. government shutdown in October 2013.

Some variance of query volume may be explained by one-off versus ongoing information needs. Users wanting to keep up with Snowden’s whereabouts or with new revelations about the NSA would repeatedly issue these queries, whereas they may have sought only once background material on privacy or surveillance. Web-browsing behavior on Wikipedia and the general Web, discussed in the following sections, provides further evidence on one-off versus ongoing information seeking.

Benchmarking against other events.

The public revelation of PRISM seems to have had much less of an influence on searches related to privacy than on other events. I contrast it with three topical issues with societal impact, chosen manually for being globally recognized events happening shortly after related media reports. During the summer of 2013, mass media reported on Hassan Rouhani’s election as president of Iran, a geopolitical issue (June 15, 2013), the U.S. Open golf tournament, a four-day major sporting event (June 13–16), and the birth of Prince

Figure 2. Changes in Web search volume for PRISM-related topics compared to other events in international politics, sports, and gossip.

Figure 3. New online documents mentioning “Snowden,” “PRISM,” and “surveillance,” respectively, over the 30 weeks following the original revelation of PRISM, June 6, 2013; the horizontal lines are averages for each seven-day period following June 6; data for October and late December 2013 are missing due to a system failure in the system used to record the data.

George of Cambridge, the “royal baby,” (July 22) (see Figure 2).

The number of queries for Rouhani increased 0.0039% on the day of his election, falling behind interest in “PRISM” but surpassing “privacy” on PRISM day. However, interest in Rouhani faded even more quickly; for example, the volume of search queries for Rouhani was significantly higher in the week following his election compared to the three weeks preceding it ($p = 0.0001$) but no more significantly in the second week ($p = 0.3$). Golf searches peaked during the four-day U.S. Open tournament (+0.25%) and were elevated in the preceding week (+0.04%) but fell sharply and below the original user base thereafter.

The birth of the royal baby showed the greatest daily increase, with +1.5% on the day of his birth, July 22, and further still, +0.41%, in the following week; the proportion of users searching for the new heir fell thereafter. Interest may be underestimated, as numbers were high before the corresponding reference period; compared to one month earlier, more than one hundred times more users searched for the royal baby on the day he was born. On 12 days during the two weeks preceding his birth, at least 0.1% of all users searched for the royal baby, and search-volume share peaked at more than 2% on July 23. Among search terms relating to PRISM, NSA and Snowden attracted their own peak search volume share with 0.4% on July 7 and June 11, respectively.

Browsing Behavior Post-PRISM

Along with Web search, browsing behavior that manifests as information seeking indicates an interest in privacy topics. I thus counted the number of users who visited webpages about Snowden, PRISM, privacy, and surveillance, using the same temporal binning as before—weeklong intervals following PRISM day. Also as before, I assessed population numbers against the three weeks prior to PRISM day.

Users browsing webpages about Snowden increased by two orders of magnitude, by far the most growth among all topics considered in my survey. Snowden stayed popular during all six weeks following PRISM day ($p < 0.00001$, t-test). PRISM and surveil-

Snowden's revelations brought few new users to privacy-enhancing technologies.

lance attracted more users as well, increasing by 95% and 250%, respectively. Whereas PRISM was able to maintain its public interest (significantly elevated page visits for all six weeks, $p < 0.0001$), the number of users visiting webpages on surveillance decreased steadily ($\rho = -0.93$, $R^2 = 0.87$) during that time and was no longer different from the reference period in the fifth week ($p = 0.22$). There was no significant increase in the number of people visiting privacy-related webpages. Numbers increased slightly in the week following PRISM day (+4%) but fell below the original levels thereafter (up to -13%).

Looking at longer-term trends, only “Snowden” continued to attract a significantly larger audience. Visits to webpages about PRISM or privacy fell to or below their original levels, and numbers for surveillance-related webpages were sporadically above the reference period in August and November/December, albeit with no clear trend.

Wikipedia is a standard online reference that fulfills general-purpose information needs. The English Wikipedia pages on PRISM and Snowden were created June 7 and 9, 2013, respectively, and data does thus not exist for the reference period. The number of page views increased significantly ($p < 0.001$) for the encyclopedia entries on privacy and surveillance, by 23% and 75%, respectively. But by week two, page views for the “privacy” Wikipedia article had already fallen to and below the reference period, with a later increase in September 2013; 25 numbers for the “surveillance” article were back to original levels by week five. Although the “Edward Snowden” article was created after the reference period, the Wikipedia article statistics provide further insight into the question of one-off versus ongoing information-seeking behavior through Web search. There is no indication Snowden would attract readership on an ongoing basis, while privacy and surveillance would not. On the contrary, the weekly interest in Snowden shrank even more drastically (-80%) than privacy and surveillance (-13%), respectively, when comparing the immediate reactions after PRISM day to longer-term evolution.

The early media coverage on PRISM reported the NSA would be “tapping directly into the central servers of nine leading U.S. Internet companies,”⁵ including Microsoft, and individuals seeking information may have consulted the privacy statements published on the corporate Website. Although consumers rarely consult privacy policies, they may have suddenly become eager to learn the details of corporate practices, including data sharing “when required by law or to respond to legal process or legal requests, including from law enforcement or other government agencies.”¹⁴ Visits increased by up to 12% in the first six weeks following PRISM day and stayed significantly above the reference period for the entire extended range ($p < 0.01$) with an upward trend ($R^2 = 0.10$, $p = 0.09$, F-test). However, I observed significantly higher numbers only in the third week following PRISM day, with no significant increase in the week immediately after the revelations ($p = 0.29$). A seasonal effect cannot be ruled out, as data from the preceding year was unavailable due to changes in how visits are counted.

Privacy-Enhancing Technology Post-PRISM

The divergence between privacy attitudes on the one hand, expressed by interest and search, and behaviors on the other, is well documented¹⁸ and was corroborated for PRISM. For example, Facebook reported diminishing trust but no impact on frequency of use of its social networking service.²⁴

Tor is a privacy-enhancing technology that allows users to conceal their location and browsing habits by routing Web traffic through multiple relays. Tor markets itself as a protective measure against network surveillance; “Browse anonymously with Tor” was featured in the *Washington Post* as the first of “five ways to stop the NSA from spying on you.”¹⁰ Tor use increased significantly— $p < 0.01$, plus up to 10%—but only in weeks three to five following PRISM day. Tor remained a niche technology; the growth of its user population was small (up by a maximum of 15,000 users in week four) and tiny compared to September 2013 when misuse of the

At the peak on July 2, 2013, “Snowden” appeared in more than 1% of all newly found online documents.

Tor infrastructure by cybercriminals drove user numbers to quadruple.

Anonymox is an alternative browser add-on service for hiding one’s IP address. Anonymox use more than doubled in 2013 to more than 200,000 users at the end of the year; I corrected for this trend in my analysis. PRISM day represented a temporary high between May and July but itself generated no spike in use ($p = 0.24$ for the largest increase in week two).

Configuration of the privacy settings in all Web browsers represents a third indicator of behavioral changes following PRISM day. Whereas the configuration of a proxy still requires technical skill, browsers are designed to be configurable by ordinary users. I used data from a sample of Internet Explorer users who consented to share usage metrics and counted how many of them selected the “privacy” tab under Internet options in the preceding month. The prevalence of this behavior increased after PRISM day by up to 2.8 percentage points compared to the reference period, approaching significance ($p = 0.02$).

Configuring the privacy options is still among the lesser-used browser features, and I found no other significant increases over the extended period. In comparing Tor/Anonymox with Internet Explorer data, I noted use of anonymizing proxies is an ongoing privacy effort, whereas adjusting browser settings toward more privacy protection is a one-off operation. The available data could still overestimate the proportion of users; for example, I counted users who opened the privacy-settings tab, without necessarily making a change, let alone activating more restrictive settings. Moreover, the sample size for Internet Explorer telemetry data was small.

Evolution of Media Coverage

I interpreted the data on consumers’ privacy behaviors, described earlier, against media coverage. Apart from a dip in the second week following PRISM day, the terms included in the survey followed a consistent upward trend. For example, more documents mentioning Snowden and surveillance appeared in week six than in week one. At the peak on July 2, 2013, “Snowden” appeared in more than

1% of all newly found online documents. And on each of the following days, “PRISM” appeared in more than 240,000 daily new documents on average while competing with other initiatives (such as Tempora and XKeyScore) (see Figure 3).

Media coverage continued over the 30 weeks following PRISM day, with no noticeable downward trend for PRISM ($\rho = -0.00002$, F-test: $p < 0.0001$), surveillance ($\rho = 0.0000$, $p = 0.001$), or Snowden ($\rho = 0.0000$, $p = 0.004$). By the end of the study period—week 30 following PRISM day—the relative daily volume of documents about Snowden was 18 times as much as on June 6, 2013.

Conclusion

This article covers the first longitudinal study of the privacy behaviors of U.S. Web users as they might have been affected by Edward Snowden's 2013 revelations about government surveillance and the general lack of communications privacy. I compared the use of privacy-enhancing technologies pre- and post-PRISM day, using Web search and browsing activity as proxies for interest in privacy and information-seeking behavior. My analysis of Web search behavior through Microsoft's Bing search engine may have introduced a bias impossible to quantify, should it exist. However, external evidence suggests Bing may be more appealing to the privacy-aware,^{19,22} meaning the small increases I observed could still represent an overestimation.

I combined high-resolution data from primary sources that indicate the new public information on PRISM led to momentarily increased interest in privacy and protection. However, the spike was much less than for other news events (such as the royal baby and the U.S. Open golf tournament). It was also less than the increased interest following the removal of privacy-enhancing functions in Facebook, Android, and Gmail.

While media coverage of PRISM and surveillance was elevated for the 30 weeks following PRISM day, many privacy behaviors faded quickly. Visits to Microsoft's corporate privacy policy page stayed high, but only certain privacy-related webpages kept

larger audiences—those on Snowden and surveillance—while Wikipedia articles about PRISM topics lost their increased readership. Snowden's revelations brought few new users to privacy-enhancing technologies; anonymizing proxies experienced increased numbers through 2013, but PRISM did not add to the momentum. While Snowden himself was the only name to consistently attract large audiences, one may argue this interest had already separated from the erosion of privacy he had revealed.

Only longitudinal studies with high temporal resolution are able to reveal the influence of privacy invasions on people's behavior. The paucity of such studies may be explained by the difficulty of obtaining expressive data. I had to rely on proxies (such as information-seeking behavior) for users' interest in privacy that cannot be observed directly. My selection of data sources was partly pragmatic, aiming for rich data providing good coverage of the same population over an extended time.

I thus opted to focus on English-language users in the U.S.; this so-called “en-us” market is a standard geographic filter for many consumer services, allowing consistent scoping across different data sets. The trade-off between narrow but good data leads to an obvious limitation, as users may have chosen to pose as “en-us” in their software settings. Still, previous internal analyses indicate their proportion is negligible. I also plan to examine the effect PRISM had (and continues to have) in various countries. My results warrant contrasting the effect of governmental versus corporate wrongdoing in privacy issues. **C**

References

1. BITKOM (Federal Association for Information Technology). Internetnutzer werden misstrauisch, July 25, 2013; http://www.bitkom.org/de/presse/8477_76831.aspx
2. Dierig, C., Fuest, B., Kaiser, T., and Wisdorff, F. Die Welt (Apr. 13, 2014); <http://www.welt.de/wirtschaft/article126882276/Deutsche-unterschaetzen-den-Wert-persoenlicher-Daten.html>
3. DuckDuckGo. DuckDuckGo Direct queries per day (28-day average), July 2014; <https://duckduckgo.com/traffic.html>
4. Dumais, S., Jeffries, R., Russell, D.M., Tang, D., and Teevan, J. Understanding user behavior through log data and analysis. Chapter in *Ways of Knowing in HCI*, Springer, New York, 2014, 349–372.
5. Gellman, B. and Poitras, L. U.S., British intelligence mining data from nine U.S. Internet companies in broad secret program. *The Washington Post* (June 7, 2013); <http://wapo.st/1888aNaq>
6. Ginsberg J. et al. Detecting influenza epidemics using

- search engine query data. *Nature* 457, 7232 (Feb. 19, 2009), 1012–1014.
7. Jackson, D., Davis, S., and Johnson, K. Obama's spy plan includes Internet: NSA's expansive reach revives an unsettled and vexing 9/11 debate. *USA Today* (June 7, 2013), A1.
8. Landau, S. Making sense from Snowden: What's significant in the NSA surveillance revelations. *IEEE Security & Privacy* 11, 4 (July–Aug. 2013), 54–63.
9. Lazer, D., Kennedy, R., King, G., and Vespignani, A. The parable of Google flu: Traps in big data analysis. *Science* 343, 6176 (Mar. 14, 2014), 1203–1205.
10. Lee, T.B. Five ways to stop the NSA from spying on you. *The Washington Post* (June 10, 2013), <http://www.washingtonpost.com/blogs/workblog/wp/2013/06/10/five-ways-to-stop-the-nsa-from-spying-on-you/>
11. Marthews, A. and Tucker, C. *Government Surveillance and Internet Search Behavior*. SSRN Working Paper, Social Science Electronic Publishing, Inc., Rochester, NY, 2014.
12. McCombs, M.E. and Shaw, D.L. The agenda-setting function of mass media. *Public Opinion Quarterly* 36, 2 (Summer 1972), 176–187.
13. Microsoft. Bing Help, 2013; <http://onlinehelp.microsoft.com/en-us/bing/ff808447.aspx>
14. Microsoft. Microsoft.com Privacy Statement, 2013; <http://www.microsoft.com/privacystatement/en-us/core/default.aspx>
15. Mozilla. anonymoX: Add-ons for Firefox, 2014; <https://addons.mozilla.org/en-US/firefox/addon/anonymox/>
16. Nakashima, E. and Markon, J. Dozens of attacks foiled, NSA says. *Washington Post* (June 13, 2013), A1.
17. Phelps, J., Gonzenbach, W., and Johnson, E. Press coverage and public perception of direct marketing and consumer privacy. *Journal of Direct Marketing* 8, 2 (Spring 1994), 9–22.
18. Preibusch, S., Kübler, D., and Beresford, A.R. Price versus privacy: An experiment into the competitive advantage of collecting less personal information. *Electronic Commerce Research* 13, 4 (Nov. 2013), 423–455.
19. Protalinski, E. Microsoft confirms Google privacy campaign to promote Bing is aimed at Apple Safari users. *The Next Web* (Sept. 20, 2012); <http://thenextweb.com/microsoft/2012/09/20/microsoft-confirms-google-privacy-campaign-aimed-apple-safari-users/>
20. Rosenblatt, S. Escaping Google's gravity: How small search engines define success. *CNET* (July 11, 2013); <http://cnet.co/15kOLCY>
21. Roznowski, J.L. A content analysis of mass media stories surrounding the consumer privacy issue, 1990–2001. *Journal of Interactive Marketing* 17, 2 (Apr. 2003), 52–69.
22. Stone, B. Facebook radically revamps its search engine. *Bloomberg Businessweek* (Jan. 15, 2015); <http://www.businessweek.com/articles/2013-01-15/facebook-radically-revamps-its-search-engine>
23. The Tor Project, Inc. *Tor Metrics Portal: Users 2013*; <https://metrics.torproject.org/users.html?graph=direct-users&country=us#direct-users>
24. Van Grove, J. Zuckerberg: Thanks NSA, now people trust Facebook even less. *CNET* (Sept. 18, 2013); http://news.cnet.com/8301-1023_3-57603561-93/zuckerberg-thanks-nsa-now-people-trust-facebook-even-less/
25. Wikipedia article traffic statistics. 2013; <http://stats.grok.se/>

Sören Preibusch (<http://preibusch.de/>) is a user experience researcher at Google, Mountain View, CA, and was at Microsoft Research, Cambridge, U.K., when this article was written.