

The Cost of Reading Privacy Policies

Aleecia M. McDonald and Lorrie Faith Cranor*

AUTHOR'S PRE-PRESS VERSION

Please cite to the published paper in:
I/S: A Journal of Law and Policy for the Information Society
2008 Privacy Year in Review issue
<http://www.is-journal.org/>

Abstract

Companies collect personally identifiable information that website visitors are not always comfortable sharing. One proposed remedy is to use economics rather than legislation to address privacy risks by creating a market place for privacy where website visitors would choose to accept or reject offers for small payments in exchange for loss of privacy. The notion of micropayments for privacy has not been realized in practice, perhaps because advertisers might be willing to pay a penny per name and IP address, yet few people would sell their contact information for only a penny.¹ In this paper we contend that the time to read privacy policies is, in and of itself, a form of payment. Instead of receiving payments to reveal information, website visitors must pay with their time to research policies in order to retain

* Aleecia M. McDonald is a PhD candidate in Engineering and Public Policy at Carnegie Mellon University. Lorrie Faith Cranor is an Associate Professor of Computer Science and of Engineering and Public Policy at Carnegie Mellon University where she is director of the CyLab Usable Privacy and Security Laboratory (CUPS).

We gratefully acknowledge Janice Tsai, Ponnurangam Kumaraguru, and Pitikorn Tengtakul for support, feedback, and helpful comments on preliminary drafts of this work. This paper benefited from insightful conversations with H. Scott Matthews, Jon M. Peha, Alessandro Acquisti, and Chriss Swaney. The authors thank Robert McGuire for his assistance with analysis. Many thanks to Michelle McGiboney and Suzy Bausch of The Nielsen Company. This research was funded in part by U.S. Army Research Office contract no. DAAD19-02-1-0389 ("Perpetually Available and Secure Information Systems") to Carnegie Mellon University's CyLab.

¹ Simson Garfinkel, *Database Nation: The Death of Privacy in the 21st Century* (Sebastopol, CA: O'Reilly & Associates, 2001), 183.

their privacy. We pose the question: if website users were to read the privacy policy for each site they visit just once a year, what would their time be worth?

Studies show privacy policies are hard to read, read infrequently, and do not support rational decision making. We calculated the average time to read privacy policies in two ways. First, we used a list of the 75 most popular websites and assumed an average reading rate of 250 words per minute to find an average reading time of 10 minutes per policy. Second, we conducted an online study of 212 participants to measure time to skim online privacy policies and respond to simple comprehension questions. We used data from Nielsen/Net Ratings to estimate the number of unique websites the average Internet user visits annually with a lower bound of 119 sites. We estimated the total number of Americans online based on Pew Internet & American Life data and Census data. Finally, we estimated the value of time as 25% of average hourly salary for leisure and twice wages for time at work. We present a range of values, and found the national opportunity cost for just the time to read policies is on the order of \$781 billion. Additional time for comparing policies between multiple sites in order to make informed decisions about privacy brings the social cost well above the market for online advertising. Given that web users also have some value for their privacy on top of the time it takes to read policies, this suggests that under the current self-regulation framework, targeted online advertising may have negative social utility.

1 Introduction

The Federal Trade Commission (“FTC”) supports industry self-regulation for online privacy.² In the late 1990s, the FTC decided that the Internet was evolving very quickly and new legislation could stifle growth. In particular, there were concerns that it was premature to legislate to protect privacy before other mechanisms evolved, especially when business was expected to offer more effective and efficient responses than FTC staff could devise. The Internet was still young, commerce on the Internet was very new, and legislators and regulators adopted a hands-off approach rather than risk stifling innovation. However, concerns remained about data privacy in general and on the Internet in particular. For example, the FTC recommended legislation to protect children’s privacy, which led to the Children’s Online Privacy Protection Act (“COPPA”) in 1998.³

Prior to COPPA, the FTC adopted Fair Information Principles (“FIPs”), a set of ideals around data use. The notion of FIPs predates the Internet; several nations adopted differing FIPs in response to concerns about credit databases on mainframes in the 1970s.⁴ While FIPs do not

² Robert Pitofsky, Chairman of the Federal Trade Commission, “Privacy Online: Fair Information Practices in the Electronic Marketplace” (prepared statement before the Senate Committee on Commerce, Science, and Transportation, Washington D.C., May 25, 2000), <http://www.ftc.gov/os/2000/05/testimonyprivacy.htm>.

³ Robert Pitofsky, Chairman of the Federal Trade Commission, “Self-Regulation and Privacy Online” (prepared statement before the Senate Committee on Commerce, Science, and Transportation, Washington D.C., July 27, 1999), <http://www.ftc.gov/os/1999/07/privacyonlinetestimony.pdf>.

⁴ Kenneth C. Laudon, “Markets and Privacy,” *Communications of the ACM* 39, no. 9 (1996): 96.

themselves carry the force of law, they provide a set of principles for legislation and government oversight. In this way they are similar to the Universal Declaration of Human Rights, in which Article 12 states the principle that “No one shall be subjected to arbitrary interference with his privacy, family, home or correspondence, nor to attacks upon his honour and reputation. Everyone has the right to the protection of the law against such interference or attacks,” but leaves the specific legal implementations of those ideals in the hands of individual nations.⁵

The five FIPs the FTC adopted in 1973 – notice/awareness, choice/consent, access/participation, integrity/security, and enforcement/redress – are a subset of the eight protections enshrined in the Organization for Economic Co-operation and Development (“OECD”) Guidelines on the Protection of Privacy and Transborder Data Flows of Personal Data.⁶ The FIP of notice underlies the notion of privacy policies, which are mechanisms for companies to disclose their practices. In 1998, the FTC commissioned a report that found while 92% of U.S. commercial websites collected some type of data, only 14% provided comprehensive notice of their practices.⁷ The FTC was concerned that the FIP of notice/awareness was not faring well on the new Internet: consumers did not know where their data went or what it might be used for.⁸

Voluntary disclosure formed the basis of an industry self-regulation approach to notice. Because privacy policies were voluntary, there were no requirements for the existence of a policy let alone any restrictions as to the format, length, readability, or content of a given privacy policy. In addition to the threat of regulatory action to spur voluntary disclosure, the FTC also used fraud and deceptive practices actions to hold companies to whatever content they did publish. In essence, while a company was not strictly required to post a policy, once published, the policy became enforceable. In one case the FTC brought action even without a privacy policy. When Cartmanager surreptitiously rented their customer lists the FTC advanced a legal theory of unfairness rather than fraud.⁹ Cartmanager provided online shopping cart software and worked with clients who promised not to sell customer data. The FTC argued that even though

⁵ UN General Assembly, *Universal Declaration of Human Rights*, art. 12, 1948, <http://www.unhcr.ch/udhr/lang/eng.pdf>.

⁶ Organisation for Economic Co-operation and Development, “OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data,” http://www.oecd.org/document/18/0,3343,en_2649_34255_1815186_1_1_1_1,00.html.

⁷ Federal Trade Commission, “Privacy Online: Fair Information Practices in the Electronic Marketplace,” 4, <http://www.ftc.gov/reports/privacy2000/privacy2000.pdf>.

⁸ *Ibid.*, 36.

⁹ David A. Stampley, *Managing Information Technology Security and Privacy Compliance* (Chicago: Neohapsis, May 2005), 3, <http://www.neohapsis.com/utility/NeoPrivacyWhitepaper.pdf> (linked to as “Privacy Compliance”).

Cartmanager did not have a privacy policy of their own to violate, they still violated the policies of their clients.¹⁰

The FTC initiated a series of studies of hundreds of commercial websites to determine how well industry self-regulation worked in what became known as Internet sweeps. Year after year, the number of companies offering privacy policies increased. By that metric it appeared the FTC was successful. However, multiple studies also showed people were reluctant to shop online because they had privacy concerns.¹¹ Recall that the FTC's charter is largely financial — barriers to new markets and commerce are a serious issue. The FTC turned to two different innovative approaches, rather than legislation or regulatory action. First, they expressed great hope for online privacy seals.¹² Two seal providers, TRUSTe and the Better Business Bureau (through BBBOnline), began certifying website privacy policies. TRUSTe requires companies to follow some basic privacy standards and document their own practices. TRUSTe also investigates consumer allegations that licensees are not abiding by their policies.¹³ However, TRUSTe has come under criticism for not requiring more rigorous privacy standards.¹⁴ In fact, one study showed that companies with TRUSTe seals typically offer less privacy-protective policies than those without TRUSTe seals.¹⁵

Second, the FTC encouraged privacy enhancing technologies (“PETs”) with the hope that PETs would put greater control directly into the hands of consumers.¹⁶ PETs include encryption, anonymity tools, and other software-based approaches. One particularly intriguing approach came from the Platform for Privacy Preferences (“P3P”) standard, which used privacy policies coded in standardized machine-readable formats. P3P user agents can determine for customers if a given website provided an acceptable privacy policy.¹⁷ Even though P3P support is

¹⁰ Federal Trade Commission, “Internet Service Provider Settles FTC Privacy Charges,” news release, March 10, 2005, <http://www.ftc.gov/opa/2005/03/cartmanager.shtm>.

¹¹ Pitofsky, “Privacy Online: Fair Information Practices in the Electronic Marketplace,” 2 (see n. 7).

¹² Pitofsky, “Self-Regulation and Privacy Online,” 5 (see n. 3).

¹³ TRUSTe, “TRUSTe Program Requirements,” <http://www.truste.org/requirements.php> (accessed Jan. 19, 2009).

¹⁴ Jamie McCarthy, “TRUSTe Decides Its Own Fate Today,” *Slashdot* (November 8, 1999), <http://slashdot.org/yro/99/11/05/1021214.shtml>.

¹⁵ Carlos Jensen and Colin Potts, “Privacy Policies Examined: Fair Warning or Fair Game?,” *GVU Technical Report 03-04* (Feb. 2003): 5, <ftp://ftp.cc.gatech.edu/pub/gvu/tr/2003/03-04.pdf>.

¹⁶ Pitofsky, “Self-Regulation and Privacy Online,” 5 (see n. 3).

¹⁷ Lorrie F. Cranor, Praveen Guduru, and Manjula Arjula, “User Interfaces for Privacy Agents,” *ACM Transactions on Computer-Human Interaction (TOCHI)* 13, no. 2 (June 2006): 135.

integrated into popular web browsers, unfortunately most users remain unfamiliar with the technology.¹⁸

1.1 *Economic Theories of Privacy Policies*

The FTC started with a set of principles, almost akin to a framework of rights, and encouraged companies to protect these rights by adopting privacy policies. Economists also see utility in privacy policies but from an entirely different basis.

Advertising economics looks at ways to turn a commodity (e.g., water) into a bundle of marketable attributes (e.g., from mountain springs). There are three types of attributes. *Search goods* are things readily evaluated in advance, for example color. *Experience goods* are only evaluated after purchase or use, for example the claims of a hair care product. *Credence attributes* cannot be determined even after use, for example nutrition content of a food. One argument for mandatory nutrition labels on food is that it converts nutrition information from a credence attribute to a search attribute: consumers can read the label prior to purchase.¹⁹ This argument applies equally well to online privacy. Without a privacy policy, consumers do not know if a company will send spam until after they have made the decision to provide their email address. With a privacy policy, consumers can check privacy protections prior to engaging in business with the site.

Another economic perspective that leads to supporting privacy policies is that since privacy is not readily observable, it cannot be properly valued by the market place. Without privacy policies, companies have all of the information about their own practices and consumers have none, leading to an information asymmetry.²⁰ Information asymmetries are one potential cause of market failure. The canonical example is of a market for used cars: sellers know if their cars are in mint condition or are lemons, but buyers may not be able to tell.²¹ Consequently, buyers need to take into account the risk of getting a bad car, and will not pay top dollar for a great car just in case they are being taken for a ride.

Privacy policies should help reduce information asymmetries because companies share information with their customers. However, researchers also note that if the cost for reading privacy policies is too high, people are unlikely to read policies. Time is one potential cost, and

¹⁸ Carlos Jensen, Colin Potts, and Christian Jensen, "Privacy practices of Internet users: Self-reports versus observed behavior," *International Journal of Human-Computer Studies* 63, no. 1–2 (2005): 212.

¹⁹ Andreas C. Drichoutis, Panagiotis Lazaridis, and Rodolfo M. Nayga, "Consumers' Use of Nutritional Labels: a Review of Research Studies and Issues," *Academy of Marketing Science Review*, no. 9 (2006): 1.

²⁰ Tony Vila, Rachel Greenstadt, and David Molnar, "Why We Can't Be Bothered to Read Privacy Policies Models of Privacy Economics as a Lemons Market," *ACM International Conference Proceeding Series* 50 (2003): 403–407.

²¹ George A. Akerlof, "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism," *Quarterly Journal of Economics* 84, no. 3 (1970): 488–500.

the time it takes to read policies may be a serious barrier.²² This approach assumes rational actors performing personal benefit-cost analysis, at least on an implicit level, to make individual decisions to read or skip privacy policies.²³ If people feel less benefit reading policies than they perceive cost of reading them, it stands to reason people will choose not to read privacy policies.

One question then is what value to place on the time it takes to read privacy policies. There is a growing literature addressing the monetary value of time, starting in the mid-1960s.²⁴ For example, urban planners estimate the value lost to traffic jams when deciding if it makes sense to invest in new roads or other infrastructure improvements.²⁵ As benefit cost analysis increased in popularity, government agencies found they had a hard time calculating economic value for “free” services like parks. One way to address their value is to estimate the time people spend traveling to parks and the value of the time they spend enjoying the parks, which again requires estimates of the value of time.²⁶ We draw upon this body of work.

In this paper we look at societal and personal opportunity costs to read privacy policies. Under the notion of industry self-regulation, consumers should visit websites, read privacy policies, and choose which websites offer the best privacy protections. In this way a market place for online privacy can evolve, and through competition and consumer pressure, companies have incentives to improve their privacy protections to a socially optimal level. In practice, industry self-regulation has fallen short of the FTC vision. First, the Internet is far more than commercial sites or a place to buy goods. While it may make sense to contrast the privacy policies of Amazon, Barnes and Noble, and O’Reilly to purchase the same book, there is no direct substitute for popular non-commercial sites like Wikipedia. Second, studies show privacy policies are hard to read,²⁷ read infrequently,²⁸ and do not support rational decision making.²⁹

²² Cranor, Guduru, and Arjula, “User Interfaces for Privacy Agents,” 135–36 (see n. 17).

²³ Alessandro Acquisti and Jens Grossklags, “Privacy and Rationality in Individual Decision Making,” *IEEE Security & Privacy* 3, no.1 (January/February 2005): 24–30.

²⁴ Gary S. Becker, “A Theory of the Allocation of Time,” *The Economic Journal* 75, no. 299 (September 1965): 493–517, <http://www.jstor.org/stable/2228949>.

²⁵ Timothy Leunig, “Time is Money: A Re-Assessment of the Passenger Social Savings from Victorian British Railways,” *The Journal of Economic History* 66 (2006): 635–73, working paper available at <http://www.lse.ac.uk/collections/economicHistory/pdf/LSTC/0905Leunig.pdf>.

²⁶ Mira G. Baron and Liliya Blekhman. “Evaluating Outdoor Recreation Parks Using TCM: On The Value of Time.” (North American Regional Science Meeting, Charleston, South Carolina, January 2002), http://ie.technion.ac.il/Home/Users/mbaron/E_21_Baron-Blekhman_Jan2_2002.pdf.

²⁷ Carlos Jensen and Colin Potts, “Privacy policies as decision-making tools: an evaluation of online privacy notices” (*Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, Vienna, Austria, April 24–29, 2004); *CHI '04 ACM* 6, no.1 (2004): 477.

Several scholars extended the FTC's vision of an implicit marketplace for privacy by examining ways to explicitly buy and sell personal information. Laudon proposed "[m]arket-based mechanisms based on individual ownership of personal information and a National Information Market (NIM) in which individuals can receive fair compensation for the use of information about themselves." Under this plan, corporations could buy "baskets of information" containing the financial, health, demographic or other data that individuals were willing to sell about themselves.³⁰ Varian sees privacy as the "right not to be annoyed" and suggests web-based contracts to sell specific information for specific uses during a fixed time frame.³¹ Yet no such market of micropayments for personal information exists. Garfinkel notes that in the current market place, where corporations re-sell information to other corporations, payments are already low. He estimates that payments to individuals for their information would be worth about a penny per name, which is far lower than most people would be willing to accept.³² Since Garfinkel's analysis, the market for personal information has been flooded with readily available information. Even stolen information is worth only about a tenth of what it used to fetch on the black market.³³ Full clickstream data sells for only 40 cents per user per month,³⁴ yet from the outrage when AOL released search term data to researchers,³⁵ it is a good guess that most people value their data at a substantially higher rate than it currently sells for on the open market. With sellers demanding more than buyers will pay, there is no zone of possible agreement, and thus it is likely that no transactions would take place.

In this paper we explore a different way of looking at privacy transactions. What if online users actually followed the self regulation vision? What would the cost be if all American Internet users took the time to read all of the privacy policies for every site they visit each year? We

²⁸ Jensen, Potts, and Jensen, "Privacy practices of Internet users: Self-reports versus observed behavior," 215 (see n. 18).

²⁹ Acquisti and Grossklags, "Privacy and Rationality in Individual Decision Making," 24-30 (see n. 23).

³⁰ Laudon, "Markets and Privacy," 99 (see n. 4).

³¹ Hal R. Varian, 1996. Economic Aspects of Personal Privacy. Faculty Working Paper, Department of Economics, Univ. of California at Berkeley, <http://people.ischool.berkeley.edu/~hal/Papers/privacy>. See sections "A simple example/search costs" and "Contracts and markets for information."

³² Garfinkel, *Database Nation*, 183 (see n. 1).

³³ Mark Trevelyan, "Stolen account prices fall as market flooded," *news.com.au*, July 15, 2008, <http://www.news.com.au/technology/story/0,25642,24023758-5014111,00.html>.

³⁴ Henry Blodget, "Complete CEO: ISPs Sell Clickstreams for \$5 a Month," *Seeking Alpha*, March 13, 2007, <http://seekingalpha.com/article/29449-competee-ceo-isps-sell-clickstreams-for-5-a-month>.

³⁵ Andrew Kantor, "AOL search data release reveals a great deal," *USA Today*, August 17, 2006, http://www.usatoday.com/tech/columnist/andrewkantor/2006-08-17-aol-data_x.htm.

model this with calculations of the time to read or skim policies, the average number of unique websites that Internet users visit each year, and the average value of time, as we present in section 2. In section 3, we combine these elements to estimate the total annual time to read policies as well as the cost to do so, both for individuals and nationwide. We discuss our findings and present our conclusions in section 4.

2 Inputs to the Model

In this section we develop a model to estimate the cost to all United States Internet users if they read the privacy policy once on each site they visit annually. We model cost both in terms of time and the economic value of that time.

We estimate the annual time to read (“ T_R ”) online privacy policies as

$$T_R = p * R * n$$

p is the population of Internet users

R is the average national reading rate

n is the average number of unique sites an Internet user visits each year

Similarly, we estimate the time to skim (“ T_S ”) online privacy policies as

$$T_S = p * S * n$$

S is the average time to skim a policy

We contrast reading to skimming because while some Internet users might read privacy policies all the way through, studies in our lab show that in practice, people may scan privacy policies for specific information they are interested in learning rather than reading policies word-for-word.³⁶

Estimating the economic value of time is more complex. As we discuss in section 2.3, based on literature in the value of time domain, leisure time is valued at a lower hourly rate than value of loss of productivity during work hours. We estimate time at home as $1/4 W$ and time at work as $2W$ where W represents average wages. Consequently we estimate not just the annual number of unique websites, but also the proportion of sites that Internet users visit at home and at work.

2.1 Time to Read or Skim Privacy Policies

We used two different methods to estimate the average time to read online privacy policies. First, we took the average word length of the most popular sites’ privacy policies and multiplied that by typical words per minute (“WPM”) reading speeds. Second, we performed an online study and measured the time it took participants to answer comprehension questions

³⁶ Robert W. Reeder, Lorrie Faith Cranor, Patrick G. Kelly, and Aleecia M. McDonald, “A User Study of the Expandable Grid Applied to P3P Privacy Policy Visualization” (*Conference on Computer and Communications Security*, Washington, D.C., October 2008); *Proceedings of the 7th ACM Workshop on Privacy in the Electronic Society (WPES '08)*, Washington, D.C., Oct. 27, 2008: 53.

about an online privacy policy. This allows us to estimate time and costs both for people who read the full policy word for word, and people who skim policies to find answers to privacy questions they have. In each case, we use a range of values for our estimates with median values as a point estimate and high and low values from the first and third quartiles.³⁷

2.1.1 Calculated Estimate to Read Popular Website Privacy Policies

We measured the word count of the 75 most popular websites based on a list of 30,000 most frequently clicked-on websites from AOL search data in October, 2005.³⁸ Because these are the most popular sites, they encompass the sorts of policies Internet users would be most likely to encounter.

As seen in Figure 1, we found a wide range of policy lengths from a low of only 144 words to a high of 7,669 words – about 15 pages of text. We used a range of word count values from the first quartile to the third quartile, with the mean value as a point estimate.

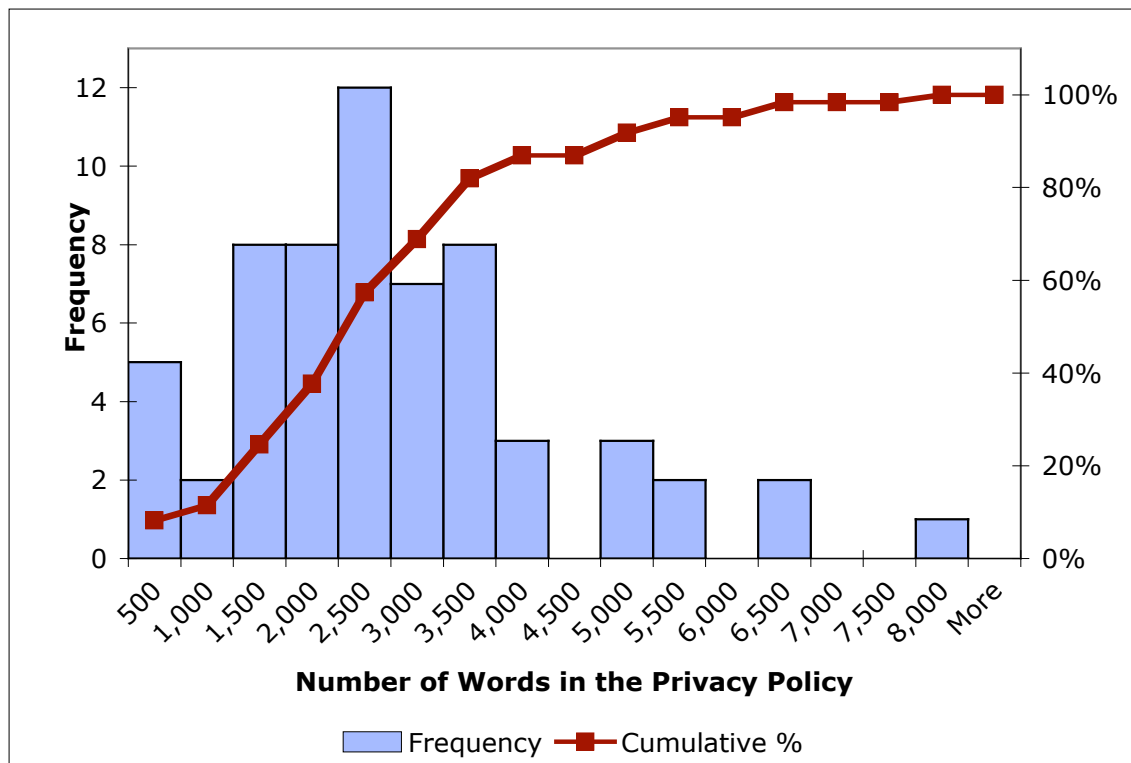


Figure 1: Probability Density Function (“PDF”) and Cumulative Distribution Function (“CDF”) of Word Counts in Popular Website Privacy Policies

³⁷ In this paper, the first quartile is the average of all data points below the median; the third quartile is the average of all data points above the median. These are single values and not a range of values. Point estimates are our single “best guess” in the face of uncertainty.

³⁸ Serge Egelman, Lorrie Faith Cranor, and Abdur Chowdhury, “An Analysis of P3P-Enabled Web Sites among Top-20 Search Results.” (Proceedings of the Eighth International Conference on Electronic Commerce, Fredericton, New Brunswick, Canada, August 14-16).

We calculated the time to read policies as the word length of common privacy policies times 250 WPM, which is a typical reading rate for people with a high school education.³⁹

Table 1: Times to read entire privacy policies for average readers

	Word Count	Reading Rate	Time to Read One Policy
Short Policy (First Quartile)	2,071 /	250 WPM =	8 minutes
Medium Policy (Median)	2,514 /	250 WPM =	10 minutes
Long Policy (Third Quartile)	3,112 /	250 WPM =	12 minutes

As seen in Table 1, we find that it takes about eight to twelve minutes to read privacy policies on the most popular sites, with a point estimate of ten minutes per policy. These estimates may be slightly low due to the jargon and advanced vocabulary in privacy policies. In addition, some people read more slowly online than on paper, which may also make these time estimates slightly low.

2.1.2 Measured Time to Skim Policies

Internet users might be more likely to skim privacy policies to find answers to their questions, or to contrast between two policies, rather than to read the policies word-for-word as envisioned in the prior section. We performed an online-study that asked participants to find the answers to questions posed about privacy protections based on the text of a privacy policy. We based our questions on concerns people have about online privacy, as studied by Cranor et al.⁴⁰ We asked five questions including “Does this policy allow Acme to put you on an email marketing list?” and “Does the website use cookies?” All answers were multiple choice, rather than short answer, so the act of answering should not have substantially increased the time to address these questions.

To ensure our results were not overly swayed by one unique policy, participants were presented with one of six different policies of varying lengths. In all, we had 212 participants from which we removed 44 outliers.⁴¹ We found that the time required to skim policies does not

³⁹ Ronald P. Carver, “Is Reading Rate Constant or Flexible?” *Reading Research Quarterly* 18, no. 2 (Winter 1983): 199, <http://www.jstor.org/stable/747517>.

⁴⁰ Cranor, Guduru, and Arjula, “User Interfaces for Privacy Agents,” 167 (see n. 17).

⁴¹ During online studies, participants are sometimes distracted by other tasks. We eliminated data points that were clearly implausible, for instance, taking 5 hours to complete a set of tasks that typically takes 20 minutes. In similar studies we have also seen responses indicative of “clicking through” the answers without reading the text. While we did have a few very speedy respondents that could mathematically be identified as outliers, we chose to retain them. For

vary linearly with length, as seen in Figure 2. We selected one very short policy (928 words), one very long policy (6,329 words) and four policies close to the typical 2,500 word length. The median times to skim one policy ranged from 18 to 26 minutes. The lowest first quartile was 12 minutes; the highest third quartile was 37 minutes. The three policies clustered near 2,500 words ranged in median times from 23 to 24 minutes and did not show statistically significant differences in mean values.⁴²

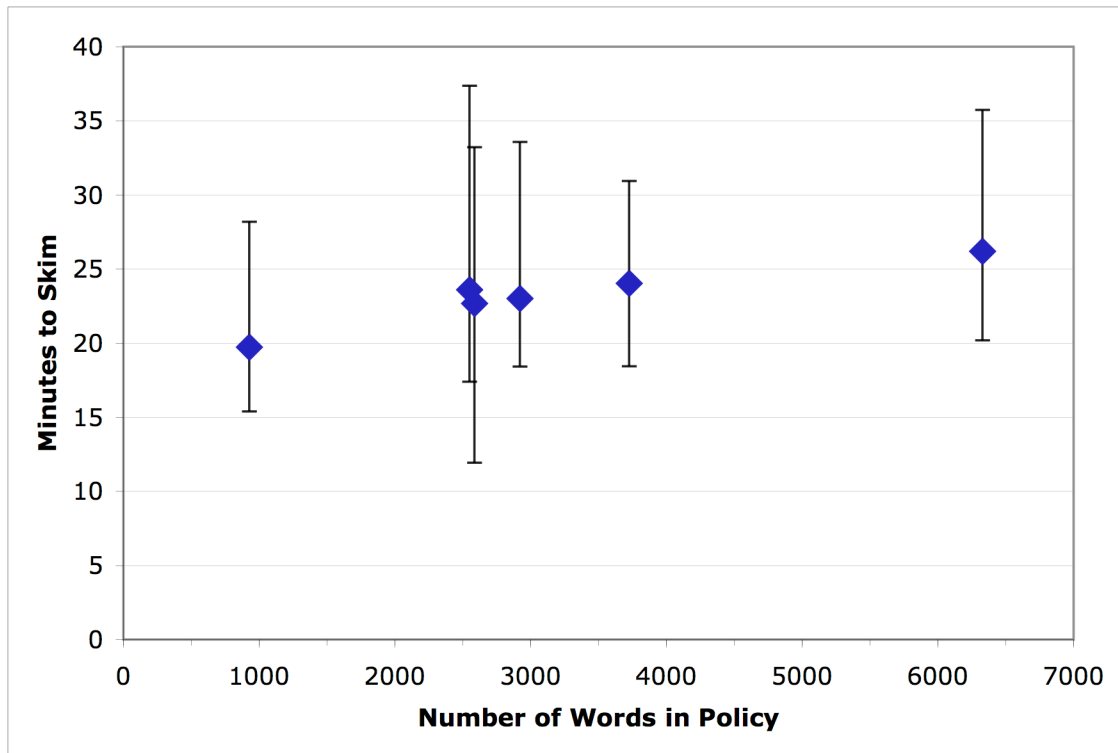


Figure 2: Median times and inter-quartile ranges to skim one privacy policy

In a prior study, we asked 93 participants to read an online privacy policy from a publishing site – the same very short 928 word policy. We asked very similar questions but included two additional questions and omitted the time to answer the first question as a training task. We found a far lower time: a point estimate of six minutes to scan a privacy policy and find relevant information. This reflects an artificially low time because, as we have since discovered, the majority of time spent answering questions is devoted to the very first question. Even though our follow up study started with a basic question, participants typically spent a third to half of their time on the very first question.

example, 3 minute response time is possibly the product of someone unusually good at the task, rather than someone who did not attempt to understand the material. In short, we favored removing and retaining outliers in ways that could slightly underestimate the times we measured.

⁴² We contrasted the 2,550 word policy to the three similar length policies using two-sided t-tests assuming unequal variance; 95% confidence interval; $p=.518, .690, .891$.

Arguably a good lower estimate of the time it takes to skim one policy is to look at the inverse of our first study: just look at the time for the first question, provided it is a question that encourages exploring the full policy. In our second study we always started with a warm up question that asked participants to identify the street address for the company and that information was always in the last few lines of the policy. Participants had to skim the full policy to answer the question. As shown in Figure 3, median times ranged from four minutes to eight minutes. The lowest first quartile of all six policies was 4 minutes; the highest third quartile was 12 minutes.

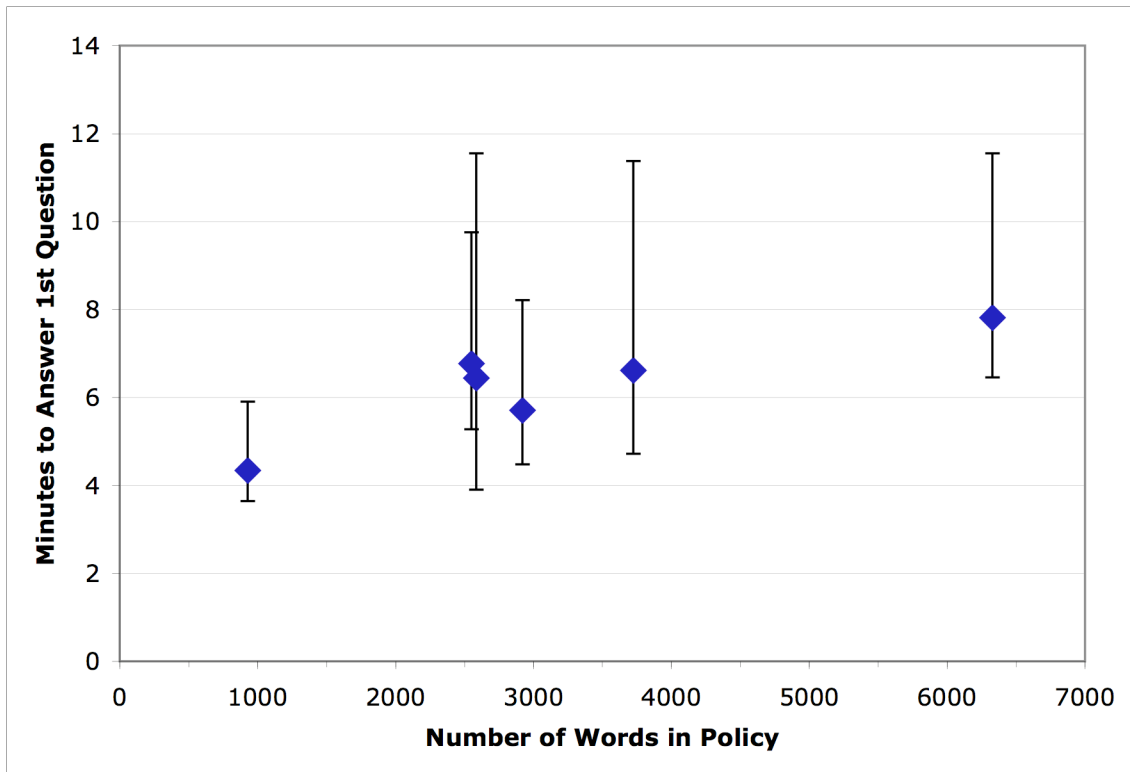


Figure 3: Median time to answer a basic question in one of six policies of different lengths, bracketed by interquartile range

One disadvantage to using just the time for the first question is that it underestimates because we only look at one question, and a very basic question at that. When asked to identify why they read privacy policies, our participants volunteered multiple interests ranging from data security, to information sales, to spam, to opt-out policies. These are captured better in the range of times reported in Figure 2. However, one advantage to using just the time for the first question is we eliminate the unsatisfying situation that we can generate longer or shorter overall time estimates just by varying the number of questions we ask.

We elected to report the more conservative estimates from just looking at the times to answer the first question, with the caveat that these numbers are lower estimates. If people were to read policies regularly, presumably they would get faster at finding information, which is another argument for a more conservative approach. We used the lowest first quartile and highest third quartile for our low and high estimates. We averaged the policies' medians as our point estimate; see Table 2.

Measured Time to Skim and Answer One Question	
Low Estimate	3.6 minutes
Point Estimate	6.3 minutes
High Estimate	11.6 minutes

Table 2: Time estimates to skim one policy and answer a basic question

2.2 Monthly Number of Unique Websites Visited

Nielsen Online reported the average number of unique websites that United States Internet users visited at home and at work during March, 2008 as 66 unique sites from work and 119 from home.⁴³ The overall average number of unique sites visited per person for the same time period was 105.⁴⁴ The overall figure is lower than the sum of sites visited from work and home because there is duplication. For example, imagine someone who visits Google both at work and at home. Google would appear once in the count of unique sites visited at work, plus once in the count of the unique sites visited at home, yet only be one unique site overall. As depicted in Figure 4, on average Internet users visit 52 different sites exclusively at work, 105 different sites exclusively at home, and 14 sites at both work and home.

⁴³ Nielsen/Net Ratings, Internet Audience Metrics, United States http://www.netratings.com/resources.jsp?section=pr_netv&nav=1 (accessed February 26, 2009) (site now updated with data reflecting the present time period).

⁴⁴ Nielsen/Net Ratings, "Nielsen Online Reports Topline U.S. Data for March 2008," news release, April 14, 2008, http://www.nielsen-online.com/pr/pr_080414.pdf.

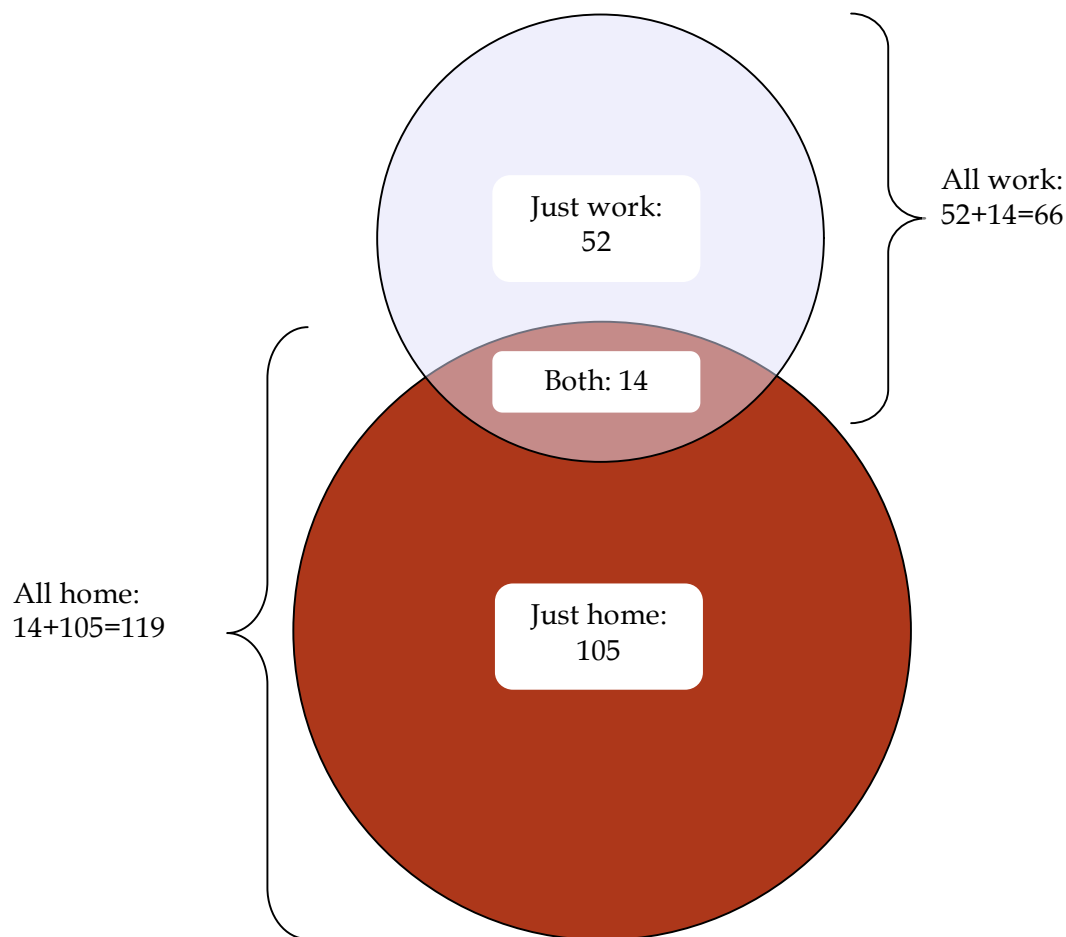


Figure 4: Locations where people read websites

We assume that if people read privacy policies, they would read them the first time they encountered a given site. We do not know where people first see the sites they visit both at work and at home. This uncertainty does not affect our time estimates but does affect our estimates for the value of that time, since time at work has a higher economic value than leisure time. As a lower bound estimate, we assume all of the sites visited at both locations are first encountered at home. As an upper bound estimate, we assume all of the sites visited at both locations are first encountered at work. For our point estimate, we split the difference and assume half are first encountered at work and half at home. These estimates are summarized in Table 3.

Table 3: Estimates of the monthly number of unique websites visited by U.S. Internet users

Estimate	Policies read at work	Policies read at home
Lower bound	52 / month	119 / month
Point estimate	59 / month	112 / month
Upper bound	66 / month	105 / month

2.3 Annual Number of Unique Websites Visited

Unfortunately, Nielsen does not collect data on the average number of websites people visit annually. They do collect weekly statistics, as shown in Table 4:

Table 4: Unique monthly and weekly websites visited by U.S. Internet users show repeat visits to many sites week after week

Location	Unique sites / month	Unique sites / week	Scale factor
Work	66	25	66%
Home	119	40	74%

People visit some of the same sites each week: if not, we would see 100 unique sites per month at home ($25 * 4$ weeks) rather than 66 (see Table 4). Ideally we would only count such sites once. From the Nielsen data we computed a scale factor, which is the percentage of sites that Internet users return to week after week. While our scale factor may not actually scale linearly over a full year it is a reasonable starting point for estimation.

We are unaware of any scholarly work that measures how many websites people visit annually. However, a 2008 study examined 25 subjects over a variable length of time and found an average of 390 unique sites during 52 to 195 days of observation.⁴⁵ The mean length of observation was 105 days. Using our point estimate of 112 unique sites per month, 390 unique sites suggests nearly all new sites each month. It seems more likely that these 25 participants, drawn from the researchers' pool of acquaintance, simply visited more sites per month than the Nielsen population. We can draw no firm conclusions. But this study does suggest, even if anecdotally, that our scale factor is not absurdly low. If anything, we may be conservative in our estimates.

For all annual estimates, we first multiplied the monthly estimate by 12 to convert from months to years, and then multiplied by the appropriate scale factor to account for visitors returning to the same sites month after month. Scale factors varied by type of estimate. As a lower bound

⁴⁵ Harald Weinreich and others, "Not quite the average: An empirical study of Web use," *ACM Transactions on the Trans Web* 2, no. 1 (February 2008): 4.

estimate for the average annual number of websites visited we multiplied by our lower observed scale factor, .66. As an upper bound annual estimate we multiplied our upper monthly estimate by our higher observed scale factor, .74. For our point estimate we used a weighted average of the observed scale factors, multiplying the monthly average work sites by the observed work scale factor of .66 and home by .74. The results are summarized in Table 5.

Table 5: Estimates of the annual number of unique websites visited by U.S. Internet users

Estimate	Scale factor		Policies read at work	+	Policies read at home	=	Total
Lower bound	.66	→	412 / year		942 / year		1354 / year
Point estimate	Weighted	→	467 / year		995 / year		1462 / year
Upper bound	.74	→	586 / year		932 / year		1518 / year

2.4 Opportunity Cost of Time

Just as the opportunity cost of time in school is a major part of the overall cost of education, Becker argued we should consider the opportunity cost of time as an implicit cost of goods and services.⁴⁶ The cost to see a play is not just the price of admission, but also the value that audience members place on their own time.⁴⁷ Economics literature suggests that time should be valued as salary plus overhead, which is the value corporations lose.⁴⁸ In the United States, overhead is estimated as twice the rate of take home pay.⁴⁹ However, that approach may not be an accurate reflection for those who work a fixed number of hours or are not in the workforce.⁵⁰ Through revealed-presences and willingness-to-pay studies, studies estimate people value their leisure time at one quarter of their take home pay.⁵¹

Taken together, this suggests that reading privacy policies at work should be valued $2W$ while reading privacy policies at home should be valued as $\frac{1}{4}W$, where W is average wages. The

⁴⁶ Gary S. Becker, "A Theory of the Allocation of Time," *The Economic Journal* 75, no. 299 (Sept. 1965): 493, <http://www.jstor.org/stable/2228949>.

⁴⁷ Ibid., 495.

⁴⁸ Timothy Leunig, "Time is Money," 493 (see n. 25).

⁴⁹ Ronald Eugen Kmetovicz, *New Product Development: Design and Analysis* (New York: Wiley-IEEE, 1992): 141.

⁵⁰ Baron and Blekman, "Evaluating Outdoor Recreation Parks Using TCM: On The Value of Time," 2 (see n. 26).

⁵¹ Leunig, "Time is Money," 638 (see n. 25).

Bureau of Labor Statistics finds an average hourly wage of \$17.93 for March, 2008.⁵² That gives us estimates of \$35.86/hour for the opportunity cost of reading privacy policies at work and \$4.48/hour for the opportunity cost of reading privacy policies at home as seen below in Table 6.

Table 6: Estimates for the value of time to read online privacy policies

Location	Average value of time
Home	\$ 4.48 / hour
Work	\$ 35.86 / hour

3 Time and Economic Value to Read Privacy Policies

In this section we use the inputs from section 2 to estimate how much time it would take for an individual to read the policies of each website she visits annually. We then use those time estimates as the basis for calculating the value of that time. In both cases we look at national figures as well as individuals.

3.1 Amount of Time to Read Privacy Policies

We multiplied the estimates for the number of unique sites American Internet users visit annually (section 2.3) by the time to read or skim privacy policies (sections 2.1.1 and 2.1.2) and by the estimated 221 million Americans online.⁵³

Table 7: Annual time estimates for reading and skimming online privacy policies

Estimate	Individual time to read	Individual time to skim	National time to read	National time to skim
Lower bound	181 hours / year	81 hours / year	39.9 billion hours / year	17.9 billion hours / year
Point Estimate	244 hours / year	154 hours / year	53.8 billion hours / year	33.9 billion hours / year
Upper bound	304 hours / year	293 hours / year	67.1 billion hours / year	64.8 billion hours / year

⁵² Bureau of Labor Statistics, Table B-3. Average hourly and weekly earnings of production and nonsupervisory workers on private nonfarm payrolls by industry sector and selected industry detail, <http://stats.bls.gov/news.release/empsit.t16.htm>.

⁵³ Nielsen/Net Ratings, "Nielsen Online Reports" (see n. 44).

We estimate that if all American Internet users were to annually read the online privacy policies word-for-word each time they visited a new site, the nation would spend about 54 billion hours reading privacy policies.

To put these figures in perspective, using the point estimate of 244 hours per year to read privacy policies per person means an average of 40 minutes a day. This is slightly more than half of the estimated 72 minutes a day people spend using the Internet.⁵⁴ This exceeds the combined percentage of Internet time devoted to shopping (1.9%) dealing with spam (6.2%) and playing games (13%) in 2005.⁵⁵ The estimated time to read privacy policies exceeds the percentage of time online that people currently spend surfing the web (45.3%).⁵⁶ One study estimates the time lost to delays in booting computers with adware as 60 hours per year per infected user, or about a quarter of the time we estimate to read privacy policies.⁵⁷ In 2000, federal income tax payers spent an estimated average of 26.4 hours completing their income taxes and nationwide, U.S. tax payers spent 3.4 billion hours completing federal income taxes⁵⁸—several times less than the amount of time we estimate for reading online privacy policies.

3.2 Value of Time to Read Privacy Policies

We multiplied the time to read or skim policies by the number of websites visited at work and the value of time at work, and added that value to the result from the same procedure for policies at home. For national costs, we again estimated 221 million Americans online.⁵⁹

⁵⁴ Norman H. Nie and others, "Ten Years After the Birth of the Internet: How Do Americans Use the Internet in Their Daily Lives?" (faculty Working Paper Stanford Institute for the Quantitative Study of Society, 2005): 4, http://www.stanford.edu/group/siqss/research/time_study_files/ProjectReport2005.pdf.

⁵⁵ *Ibid.*, 6.

⁵⁶ *Ibid.*, 5.

⁵⁷ Roger Thompson, "Minimizing liability and productivity risks: How to control the impacts of spyware, hacker tools and other harmful applications," *Computer Associates*, Oct. 2004, <http://www.ameinfo.com/pdfdocs/51515.pdf>.

⁵⁸ John L. Guyton, Adam K. Korobow, Peter S. Lee, and Eric J. Toder, "The Effects of Tax Software and Paid Preparers on Compliance Costs," *National Tax Journal* 58, no. 3 (2005): 441.

⁵⁹ Nielsen/Net Ratings, "Nielsen Online Reports" (see n. 44).

Estimate	Individual cost to read	Individual cost to skim	National cost to read	National cost to skim
Lower bound	\$2,533 / year (work: \$1,970; home: \$563)	\$1,140 / year (work: \$886; home: \$253)	\$559.7 billion / year (work: \$435 B; home: \$124 B)	\$251.9 billion / year (work: \$196 B; home: \$56 B)
Point	\$3,534 / year (work: \$2,791; home: \$743)	\$2,226 / year (work: \$1,758; home: \$468)	\$781 billion / year (work: \$617 B; home: \$164 B)	\$492 billion / year (work: \$389 B; home: \$103 B)
Upper bound	\$5,038 / year (work: \$4,203; home: \$835)	\$4,870 / year (work: \$4,063; home: \$807)	\$1.1 trillion / year (work: \$929 B; home: \$184 B)	\$1.1 trillion / year (work: \$898 B; home: \$178 B)

We estimate that if all American Internet users were to annually read online privacy policies word-for-word each time they visited a new site, the nation would lose the value of about \$781 billion from the opportunity cost value of the time to read privacy policies.

Again, to put this in perspective, in 2005 the average cost to connect to the Internet was \$237/year for dial up and \$508/year for high speed access.⁶⁰ This suggests the value of time lost to reading privacy policies would eclipse the cost of high speed Internet access, several times over. In 2007, United States online sales were approximately \$260 billion⁶¹—more than the cost to businesses if their employees were to read privacy policies on corporate time.

4 Discussion and Conclusions

We estimate that reading privacy policies carries costs in time of approximately 201 hours a year, worth about \$3,534 annually per American Internet user. Nationally, if Americans were to read online privacy policies word-for-word, we estimate the value of time lost as about \$781 billion annually.

These estimates presume that people visit sites, read the policies once a year, and then carry on their business as before. Yet the FTC vision of self-regulation presumes that, at least for consumer sites, Internet users will visit multiple sites to comparison shop for acceptable privacy

⁶⁰ Scott J. Savage and Donald Waldman, "Broadband Internet access, awareness, and use: Analysis of United States household data," *Telecommunications Policy* 29, no. 8 (2005): 615.

⁶¹ "Online sales spike 19 percent," *CNN Money.com* (May 14, 2007), http://money.cnn.com/2007/05/14/news/economy/online_retailing/index.htm.

practices. The true cost of adherence to the self-regulation vision is perhaps on the order of double the costs we estimate, depending on which percentage of sites have ready substitutes and how many sites people are expected to compare. True costs also include Internet connectivity fees, which we did not attempt to quantify.

In the opposite direction, media consolidation means that multiple sites may share one privacy policy. While consolidation itself poses increased threats to online privacy, in some cases it may actually reduce the cost of reading privacy policies because there are fewer unique policies to read. We do note that the resulting privacy policy when companies merge may be more complex and longer than either of the individual policies. Another issue is that people may not care about all possible privacy threats. For instance, if they only care about credit card theft, and they visit a site that does not collect credit card numbers, they may not feel the need to protect any information. Thus, arguably, they do not need to read the policy at every site they visit, but only a subset of sites.

The value of all online advertising in the United States was about \$21 billion in 2007.⁶² Many, though by no means all, online privacy concerns stem from advertisers amassing information about Internet users in order to present ads targeted to specific demographics. The current policy decisions surrounding online privacy suggest that Internet users should give up an estimated \$781 billion of their time to protect themselves from an industry worth substantially less. This is not to say online advertising should be banned. Sales from direct mail are approximately an order of magnitude higher than advertising costs and the cost of online advertisements similarly understates the full market.⁶³ But it appears the balance between the costs borne by Internet users versus the benefits of targeted ads for industry is out of kilter, at least as envisioned by the FTC's solution that Internet users read privacy policies.

Some Internet users may realize a benefit from targeted advertisements; for example Amazon's ability to suggest additional books they might enjoy based on prior purchase history. Yet on the whole, advertisements are usually seen as an economic "bad" rather than a "good" because participants would pay money to eliminate ads from most types of media.⁶⁴ While an analysis of the net social welfare changes created by online advertisement is beyond the scope of this paper, we do suggest that any such cost-benefit analysis should include the value of time for reading privacy policies.

Preliminary work from a small pilot study in our laboratory revealed that some Internet users believe their only serious risk online is they may lose up to \$50 if their credit card information is

⁶² Interactive Advertising Bureau, "Internet Advertising Revenues Top \$21 Billion in '07, Reaching Record High," news release, May 15, 2008, http://www.iab.net/about_the_iab/recent_press_releases/press_release_archive/press_release/299609.

⁶³ Ross D. Petty, "Marketing without consent: Consumer choice and costs, privacy, and public policy," *Journal of Public Policy and Marketing* 19, no. 1 (Spring 2000): 45.

⁶⁴ Gary S. Becker and Kevin M. Murphy, "A Simple Theory of Advertising as a Good or Bad," *The Quarterly Journal of Economics* 108, no. 4 (Nov. 1993): 961, <http://www.jstor.org/stable/2118455>.

stolen. For people who think that is their primary risk, our point estimates show the value of their time to read policies far exceeds this risk. Even for our lower bound estimates of the value of time, it is not worth reading privacy policies. This leads to two implications. First, seeing their only risk as credit card fraud suggests Internet users likely do not understand the risks to their privacy. As an FTC report recently stated, “it is unclear whether consumers even understand that their information is being collected, aggregated, and used to deliver advertising.”⁶⁵ Second, if the privacy community can find ways to reduce the time cost of reading policies, it may be easier to convince Internet users to do so. For example, if we can help people move from needing to read policies word-for-word and only skim policies by providing useful headings, or if we can offer ways to hide all but relevant information – and thus reduce the effective length of the policies – more people may be willing to read them.

The privacy community and industry groups have responded with several attempts to improve privacy policies. Layered privacy notices specify a few high-level and standardized topics for a one-screen summary of the policy, then link to the full privacy policy for more information.⁶⁶ The Platform for Privacy Preferences (“P3P”) is an XML-based specification that enables policy authors to code privacy policies in machine-readable format⁶⁷ which fosters comparison between policies in a standardized way, and provides a common format for user agents to help Internet users find acceptable policies. Privacy Bird is a web browser add-on that uses P3P to generate a short privacy report that presents information in bulleted lists with sections that expand and contract to show and hide sections of the privacy policy.⁶⁸ The P3P Expandable Grid is also built on P3P and uses icons to convey what information companies collect and how they use it.⁶⁹ Icons in the Privacy Finder search engine convey how well a given P3P policy matches users’ preferences. A Privacy Finder user study demonstrated that Internet users will pay a premium for products from sites rated as more privacy protective.⁷⁰ Both education and enhanced privacy policy formats may help Internet users gain the tools they need to protect themselves online.

⁶⁵ Federal Trade Commission, “Protecting Consumers in the Next Tech-ade: A Report by the Staff of the Federal Trade Commission,” March 2008, 11, <http://www.ftc.gov/os/2008/03/P064101tech.pdf>.

⁶⁶ Center for Information Policy Leadership, “Ten steps to develop a multilayered privacy policy,” 2007, http://www.hunton.com/files/tbl_s47details%5Cfileupload265%5C1405%5Cten_steps_whitepaper.pdf.

⁶⁷ W3C Working Group, “The Platform for Privacy Preferences 1.1 (P3P1.1) Specification,” November 2006, <http://www.w3.org/TR/P3P11>.

⁶⁸ Cranor, Guduru, and Arjula, “User Interfaces for Privacy Agents,” 149 (see n. 17).

⁶⁹ Reeder, Cranor, Kelley, and McDonald, “A User Study of the Expandable Grid Applied to P3P Privacy Policy Visualization,” 9 (see n. 36)

⁷⁰ Janice Tsai, Serge Egelman, Lorrie Cranor, and Alessandro Acquisti, “The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study” (*Workshop on the Economics of Information Security*, Pittsburgh, PA, June 7–8, 2007), 15, <http://weis2007.econinfosec.org/papers/57.pdf>.

Finally, some corporations take the view that their users should read privacy policies and if they fail to do so, it is evidence of lack of concern about privacy. Instead, we counter that websites need to do a better job of conveying their practices in useable ways, which includes reducing the time it takes to read policies. If corporations cannot do so, regulation may be necessary to provide basic privacy protections. Disclosure legislation may be insufficient: adding more text to policies that most consumers do not read does increase transparency, but may otherwise be of limited practical utility.