Privacy Papers for Policy Makers

2009-2010

The publication of "Privacy Papers for Policy Makers" was supported
by AT&T, LexisNexis, Microsoft and Procter & Gamble

September 15, 2010

We are delighted to provide you with this Journal featuring works selected by the Future of Privacy Forum Advisory Board as the best "Privacy Papers for Policy Makers," representing cutting-edge research and analytical work on a variety of privacy topics.

We solicited papers that clearly analyzed current and emerging privacy issues, and either proposed achievable short-term solutions or offered fresh analysis that could lead to new approaches and solutions. Academics, privacy advocates and Chief Privacy Officers on FPF's Advisory Board reviewed all submitted papers, emphasizing clarity, practicality and overall utility as the most important criteria for inclusion.

We hope this relevant and timely scholarship helps inform policy makers in Congress, at the FTC, and in other federal and state agencies as they address privacy issues. This compilation is also being provided to policy makers abroad.

We want to thank AT&T, LexisNexis, Microsoft and Procter & Gamble for their special support of this project. And thank you for your interest in exploring new ways to think about privacy.

Sincerely yours,

Christopher Wolf
Founder and Co-chair

Jules Polonetsky
Director and Co-chair

Future of Privacy Forum Advisory Board

Annie I. Antón
North Carolina State University

Elise Berkower
The Nielsen Company

Joan (Jodie) Z. Bernstein
Counsel, Kelley Drye & Warren, LLP and former director of the Bureau of Consumer Protection at the Federal Trade Commission

Bruce Boyden
Marquette University Law School

Kathryn C. Brown
Verizon

James Byrne
Lockheed Martin

Ryan Calo
Center for Internet & Society, Stanford Law School

Dr. Ann Cavoukian
Ontario Privacy Commissioner

Danielle Citron
University of Maryland Law School

Lorrie Faith Cranor
Carnegie Mellon University

Mary Culnan
Bentley University

Simon Davies
Privacy International

Michelle Dennedy
Cloud Computing, Sun Microsystems

Carol DiBattiste
LexisNexis

Benjamin Edelman
Harvard Business School

Scott Goss
Qualcomm

Leslie Harris
Center for Democracy and Technology (CDT)

David Hoffman
Intel

Marcia Hoffman
Electronic Frontier Foundation

Andy Holleman
Qwest Communications

Chris Hoofnagle
Berkeley Center for Law & Technology

Pamela Jones Harbour
Fulbright & Jaworski, Former FTC Commissioner

Sandra Hughes
Procter & Gamble

Nuala O'Connor Kelly
General Electric

Ian Kerr
University of Ottawa, Faculty of Law

Brian Knapp
Loopt

Gerard Lewis
Comcast

Brendon Lynch
Microsoft

Fran Maier
TRUSTe

Terry McQuay
Nymity

Rena Mears
Deloitte & Touche LLP

Scott Meyer
Better Advertising

Doug Miller
AOL

Paul Ohm
University of Colorado Law School

Adam Palmer
Symantec

Harriet Pearson
IBM

Robert Quinn
AT&T

MeMe Rasmussen
Adobe

Russell Schrader
Visa

Paul Schwartz
University of California-Berkeley School of Law

Scott Shipman
eBay

Daniel Solove
George Washington University Law School

Zoe Strickland
Walmart

Tim Sparapani
Facebook

Omar Tawakol
BlueKai

Omer Tene
College of Management School of Law, Rishon Le Zion, Israel

Anne Toth
Yahoo!

Steven Vine
Datran Media

Michael Zimmer
University of Wisconsin-Milwaukee

Table of Contents

Privacy on the Books and on the Ground

Kenneth A. Bamberger* and Deirdre K. Mulligan**

(Forthcoming, Stanford Law Review, vol. 63 (2011))

Full paper available at: www.futureofprivacy.org/the-privacy-papers/

Executive Summary

U.S. privacy law is under attack. Scholars and advocates criticize it as weak, incomplete and confusing, and argue that it fails to empower individuals to control the use of their personal information. The most recent detailed inquiry into corporate treatment of privacy, conducted in 1994, frames these critiques, finding that firms neglected the issue in their data management practices because of the ambiguity in privacy mandates and lax enforcement. As Congress and the Obama Administration consider privacy reform, they encounter a drumbeat of arguments favoring the elimination of legal ambiguity by adoption of omnibus privacy statutes, the EU's approach.

These critiques may present an accurate description of privacy law "on the books." But the debate has strangely ignored privacy "on the ground"—since 1994, no one has conducted a sustained inquiry into how corporations actually manage privacy, and what motivates them. This omission is especially striking because the neglect of the '90s has been replaced by a massive dedication of corporate resources to privacy management, the inclusion of privacy officers at the c-suite level, and the employment of a 6,500-strong cadre of privacy professionals.

This Article presents findings from the first study of corporate privacy management in fifteen years, involving qualitative interviews with Chief Privacy Officers identified by their peers as industry leaders. Spurred by these findings, we present a descriptive account of privacy "on the ground" that upends the terms of the prevailing policy debate, and offers four important insights for policymakers considering reform.

Our Study of Corporate Privacy Officers

Our alternative account identifies elements neglected by the traditional story: the emergence of the Federal Trade Commission (FTC) as a privacy regulator, the increasing influence of privacy advocates, market and media pressures for privacy protection, and the rise of privacy professionals. We trace the ways in which these players supplement a privacy debate largely focused on processes, such as notice and consent mechanisms, with a growing corporate emphasis on substance: preventing violations of consumers' expectations of privacy.

Two alterations to the legal landscape contribute to this definitional shift. First, the substantive definition tracks the emergence of the FTC as a roving regulator with broad yet ambiguous power to evaluate privacy practices in the marketplace through its consumer protection lens. The FTC's mandate to protect consumers from "unfairness" and "deception" permits dynamic regulation that evolves with changing contexts, and forces corporate practices to develop accordingly. Second, state security breach notification laws raised the soft and hard costs of mismanaging personal information. Together these changes led identified industry leaders to integrate substantive considerations of consumers' privacy expectations into their workflows, rather than leaving privacy to the lawyers and their process-based "click through if you 'consent' to the privacy policy" approach.

* Assistant Professor of Law, University of California, Berkeley, School of Law (Boalt Hall).

** Assistant Professor, University of California, Berkeley, School of Information.

Policy Implications

Our grounded account should inform privacy reforms. Specifically, it provides four important insights for policymakers considering privacy initiatives.

First, our account supports the argument that calls for federal regulation structured exclusively around fair information practice principles (FIPPS) are ill-advised. Our interviews indicated ways that FIPPS's focus on procedural mechanisms intended to give individuals "control" over their information—rather than on substantive limits on what corporations can do with that information—was insufficient to guide corporate behavior, particularly in times of profound technical or market change. Indeed, they reported, the focus on procedure rather than outcomes can create stumbling blocks for CPOs by positioning them once again as the "no" person. Thus many of our interviewees discussed efforts to transform internal perceptions about privacy from a compliance-oriented, rule-dominated legal hurdle to cross at the end stage of product design, to a consultation and dialogue about how technical designs, business strategies and policies can respect consumers' expectations and support trust in their companies from the beginning. Our interviewees further suggested that, without a substantive touchstone, a data-protection regime can focus resources on developing a host of often meaningless consent processes, which must be designed and redesigned in an effort to do better—where the meaning of "better" is unclear. They further predicted that the limitations of consent as the dominant fallback for protecting consumer privacy would be exacerbated by the increasing trend toward networks, embedded devices and increasingly personalized services.

While FIPPS remain an important touchstone for information privacy in the U.S., they should not be the exclusive criterion for regulatory reforms. FTC enforcement aimed at protecting consumers' reliance on conventional information flows has brought greater substance and meaning to an area routinely critiqued for its formalism. Protecting existing social norms about information use, rather than leaving each individual to the mercy of the marketplace, is key to addressing both collective and individual interests, for in the words of scholar Julie Cohen, while "[p]rivacy self-defense operates at the individual level . . . surveillance operates at the collective level;" thus "the logics of surveillance require a considered, collective response."

Second, our account identifies the important role that the FTC plays as a forum for shaping a collective understanding of privacy among advocates, industry, academics and regulators. While the FTC's function as roving enforcement agency has been especially significant, its threat of coercive authority leverages an even deeper role in developing a cross-field understanding of privacy through workshops, fact-finding investigations, and other soft-law techniques to flesh out the meaning of its ambiguous privacy mandate.

The collective engagement prompted by these regulatory choices has yielded groundbreaking outcomes. The FTC's combination of enforcement threats, with its centrality in creating a social network of privacy advocates, offers a model for avoiding both the shortcomings of static top-down command-and-control regulatory approaches and the ways in which reliance on bottom-up self-regulation alone can allow private interests to subvert public goals. Moving forward, this model should guide the choice and design of whatever regulatory institutions take the lead on information privacy in the corporate sector. They must both possess and use regulatory tools that exploit market, corporate and advocacy capacity to develop collective understanding of risks and solutions to future privacy problems.

Julie E. Cohen, Privacy, Visibility, Transparency, and Exposure. 75 U. CHI L. REV. 1 (2008).

Third, our account begins to illuminate the ways in which corporate privacy professionals impart meaning and structure to societal privacy concerns within corporations. While no proposal for a dedicated U.S. privacy agency has garnered public or political support, a cadre of privacy specialists has developed instead within the private sector among companies, advocacy organizations and academia. In the absence of a DPA staffed with data protection experts, and faced with increasing ambiguity as to what privacy requires, corporations depend on these new professionals to guide them through the challenges wrought by evolutions in technology and business practice. These professionals do not view themselves as compliance officers, but rather as norm entrepreneurs. Empowered by external threats that support their entrepreneurial efforts, they offer a unique capacity to embed privacy—as trust and consumer expectations—into the corporate psyche as well as business operations.

Choices about regulatory form will affect the ability to leverage these professionals, e.g., to empower them within their own organizations in ways that pushes privacy further into corporate culture. A decision to redirect privacy regulation towards more rule-bound governance, for example, might diminish the need for corporations to rely on high-level internal advocates of privacy concerns. As society becomes more pervasively networked, and privacy protection requires ongoing and on-the-ground attention to dynamic privacy interests that manifest in very different ways within different firms, institutional reforms should be attentive to preserving the benefits flowing from this embedded class of professionals, and seek to empower rather than displace them.

Finally, as the privacy community reflects upon the key global instruments of data protection, our account underscores the importance of empirical inquiry and thick institutional engagement in considering contested issues of regulatory strategy, technological complexity, social and institutional networks, and the protection of individual and communal interests in the private sphere. If privacy is to be protected in an increasingly connected world, debates over its formal regulation must increasingly be informed by the ways that today's frameworks operate on the ground.

* * *

In sum, while we do not quibble with efforts to expand procedural mechanisms to empower individuals to control their personal information, doing so in a way that (1) eclipses robust, substantive definitions of privacy and the protections they are beginning to produce; (2) constrains the regulatory flexibility that permits their evolution; or (3) disempowers corporate privacy professionals would destroy important tools for limiting corporate over-reaching, curbing consumer manipulation, and protecting shared expectations about the personal sphere on the Internet, and in the marketplace.

Authors

Kenneth A. Bamberger teaches and writes about Administrative and Constitutional Law, and Business Regulation, at the University of California, Berkeley, School of Law. His research focuses on issues of technology in governance, regulatory decision making and corporate compliance. Bamberger is affiliated with the Center for the Study of Law and Society, and with the Berkeley Center for Law, Business and the Economy.

Deirdre K. Mulligan teaches and writes on law and policy issues related to information and communication technology. Her research focuses on issues of information privacy, information security, surveillance and the interplay between legal rules and technical systems. A professor at the UC Berkeley School of Information, Mulligan is a faculty director of the Berkeley Center for Law and Technology, a board member of the Center for Democracy & Technology and Co-chair of Microsoft's Trusted Computing Academic Advisory Board.

What is Privacy Worth?

Alessandro Acquisti, Leslie John and George Loewenstein[2]

Full paper available at: www.futureofprivacy.org/the-privacy-papers/

Executive Summary

Understanding the value that individuals assign to the protection of their personal data is of great importance not only to researchers and businesses, but to policy makers who are often required to choose between policies that trade off privacy against other desirable goals. In recent years there has been no shortage of empirical studies attempting to quantify individual privacy valuations in diverse contexts. Some of these studies-- as well as anecdotal evidence about the growing popularity of blogs, online social networks and other information-sharing social media-- suggest that even ostensibly privacy conscious individuals are likely to share sensitive information with strangers. Applying the economics principle of "revealed preferences," some have concluded that our society, quite simply, does not place much value on privacy. Is it really possible, however, to measure the value that people place on privacy? And has "less privacy" truly become the new social norm, as a prominent Web 2.0 CEO has recently claimed?

In this manuscript, we challenge the view that true privacy valuations can be precisely estimated, and argue that the revealed preferences argument does not necessarily support the conclusion that people, on average, do not care for privacy. We undertook a series of experiments to understand privacy valuations (and privacy decision making) through the lenses of behavioral economics and decision research. We found that privacy valuations are highly vulnerable to subtle, non-normative influences. Specifically, we found evidence of order and endowment effects and non-normal distributions of valuations. In particular, we found that individuals assign markedly different values to the privacy of their data depending on the order in which they consider different offers for that data, or whether they consider the amount of money they would accept to disclose otherwise private information, or the amount of money they would pay to protect otherwise public information. We found the gap between such values to be larger than what is usually estimated in comparable studies of other private goods. In addition, we found evidence that privacy valuations are not normally or uniformly distributed, but U-shaped, and clustered around extreme, focal values.

These results paint a more nuanced and detailed picture of privacy valuations than the one currently in the literature. They suggest that the value of privacy, while not entirely arbitrary, is highly malleable and sensitive to non-normative factors. These results stand in contrast to other proposed estimates of privacy valuations, as well as to the view that an economic revealed preferences argument (such as consumers' eagerness to disseminate personal information to friends and strangers alike, or their relative disinterest in freely available protective technologies, such as Tor or PGP) should support the conclusion that consumers do not care for privacy. In one of our experiments, subjects were *five times more likely to reject cash offers for their data* if they believed that their privacy would be, by default, protected, than if they didn't enjoy such belief.

The policy, managerial and research implications of these findings are significant. Individuals' marketplace decisions about their data are sometimes taken as representing their true and final preferences towards protection or revelation of personal data, and therefore become an instrument for the assignment of societal resources to privacy issues. For example, the observation that individuals give away their personal

[2]acquisti@andrew.cmu.edu, lkjohn@andrew.cmu.edu, and gl20@andrew.cmu.edu.

information for small rewards has permeated the policy debate and has been used to argue against privacy regulation, on the grounds that if consumers wanted more privacy they would, in fact, ask for it and take advantage of opportunities for its protection. However, as our experiments demonstrate, such "revealed preferences" arguments should not alone justify the uncritical conclusion that even privacy conscious consumers will always be unlikely to pay for online privacy. If individual decisions regarding privacy are malleable to non-normative factors, then such arguments lose their normative standing. The answers to questions such as "What is privacy worth?" and "Do people really care for privacy?" depend not just on whom, but *how* we ask: in our experiments. Subjects who started from positions of greater privacy protection were five times more likely than other subjects to forego money to preserve that protection. Combined with the difficulty of making the "right" privacy decisions for consumers, such findings suggest that market outcomes alone may not necessarily tell us the final and last word on consumer data protection.

The peer-reviewed conference version of this paper was accepted in: *The Twentieth Workshop on Information Systems and Economics (WISE)*, December 2009, Arizona Biltmore Resort & Spa, Phoenix, AZ, http://pages.stern.nyu.edu/~bakos/wise/.

The Journal version is under review with: *Information Systems Research*

Authors

Alessandro Acquisti is an Associate Professor of Information Technology and Public Policy at the Heinz College, Carnegie Mellon University. He is also a member of Carnegie Mellon CyLab, the CMU Usable Privacy and Security Laboratory, and a fellow of the Ponemon Institute. His research focuses on the economics and behavioral economics of privacy and information security and privacy in online social networks. A recipient of the PET *Award for Outstanding Research in Privacy Enhancing Technologies* and the IBM Best Academic Privacy Faculty Award, he has also been awarded research grants from the National Science Foundation, TransCoop Foundation, Google and Microsoft.

His manuscripts have been published in leading journals of multiple disciplines including *Proceedings of the National Academy of Science, Journal of Consumer Research, Marketing Science, Information Systems Research, Journal of Comparative Economics* and *ACM Transactions on the Internet*. He has edited books and book chapters, led international conference proceedings and delivered international keynotes. His research has been featured on NPR and CNN and in *The New York Times*, *The Wall Street Journal* and *The Washington Post*. His 2009 study on the predictability of Social Security numbers was featured in the "Year in Ideas" issue of *The New York Times Magazine*. Acquisti holds a Ph.D. from UC Berkeley and Masters degrees from UC Berkeley, the London School of Economics and Trinity College.

A doctoral candidate in Behavioral Decision Research at Carnegie Mellon University, **Leslie John** conducts research at the intersection of psychology, economics and public policy. Her research focuses primarily on understanding when and why people are willing to divulge sensitive personal information, and she is also interested in developing incentive schemes informed by behavioral economics to help people improve their own health. Her research has been published in many journals from the *Journal of Consumer Research* to the *Journal of the American Medical Association*. John has a Bachelor's degree in Psychology from the University of Waterloo and a Master's degree in Behavioral Decision Research from CMU.

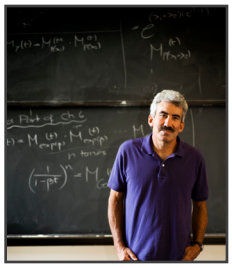

The Herbert A. Simon Professor of Economics and Psychology at Carnegie Mellon University, **George Loewenstein** received his Ph.D. from Yale University. He has held academic positions at the University of Chicago and Carnegie Mellon University, and fellowships at the Center for Advanced Study in the Behavioral Sciences, the Institute for Advanced Study in Princeton, the Russell Sage Foundation and the Institute for Advanced Study in Berlin. A co-founder of the fields of behavioral economics and neuroeconomics, Loewenstein is a past president of the Society for Judgment and Decision Making and a fellow of the American Academy of Arts and Sciences.

His research focuses on applications of psychology to economics, with an emphasis on decision making over time, bargaining and negotiations, psychology and health, law and economics, the psychology of adaptation, the role of emotion in decision making, the psychology of curiosity, conflict of interest, and "out of control" behaviors such as impulsive violent crime and drug addiction. He has published more than 100 journal articles, numerous book chapters, and has written or edited six books on topics ranging from intertemporal choice to behavioral economics to emotions.

Misplaced Confidences: Privacy and the Control Paradox

Laura Brandimarte, Alessandro Acquisti and George Loewenstein

Full paper available at: www.futureofprivacy.org/the-privacy-papers/

Executive Summary

Modern information technologies grant us great power to broadcast our personal information to the world, but afford us much less control over how that information, once disseminated, will be used. The contrast between control over the *publication*, and control over the access and *usage* of personal information, was at the center of our enquiry. People display specific concerns about the way in which their private information may be accessed and used by strangers (*Consumer Reports* poll, 2008; Norberg et al., 2007). However, at the moment of deciding whether to reveal personal information, the consideration that the information provided may later become available to individuals other than the intended recipients may remain as "latent" as such later access appears distant. The individual, therefore, may not feel concerned about the relative lack of control on future access to and usage of their data; such concerns may be trumped by the satisfactory perception of having control over the very act of revealing, and publishing, the information.

Control over publication makes the lack of control over usage less salient. Our manuscript tested such conjecture: we hypothesize that one of the psychological mechanisms that leads people to expose themselves to such a large extent is a control paradox on the information they reveal. Since we have control over the publication of our private information, we give less importance to control (or lack thereof) of the accessibility and use of that information by others. In order to investigate issues of control in privacy decision making, we ran three survey-based experiments and administered them to students at a North American University. Across all experiments, we manipulated the subjects' control over information publication, without altering the actual conditions of access to, and usage of, the information they were asked to reveal. In two experiments we decreased subjects' control, relative to a baseline condition; in one experiment we increased it. We measured subjects' propensity to reveal sensitive information about themselves as a function of the amount of control they felt over the publication of their responses. Our results indicate that people may suffer from what we call a control paradox on personal information: more control over the publication of private information makes control over information access and use by others appear less salient, which consequently both decreases individuals' privacy concerns and increases their willingness to publish sensitive information about themselves. Vice versa, individuals with less control over the publication of their private information may face increased privacy concerns, exhibiting lower willingness to publish sensitive information.

These results are significant on two levels. On a theoretical level, they challenge the traditional scholarly construct of privacy as "control" of personal information flows. Normatively, we have no doubt that granting individuals control on how their personal information is disseminated and used is an important (albeit not necessarily sufficient) condition for "privacy" protection. Positively, however, our results indicate that actually granting users control over their data is not guaranteed to make it easier for them to achieve some desired abstract balance between information revelation and information protection; if anything, the ultimate effect seems to be, paradoxically, to induce them to reveal more personal information, even when this may expose them to larger risks.

On a practical level, our results challenge the view that Internet operators can soothe privacy concerns by simply affording users more control over their data. The Internet in general, and Web 2.0 applications in particular, are granting individuals vast powers to disseminate personal thoughts and information to others. One of the Web 2.0 entrepreneurs' responses to the privacy concerns raised by their technologies has been the observation that such technologies also grant great user control in terms of to whom, when and how to present one's online persona. For instance, in a 2004 interview, then Tribe.net CEO Mark Pincus claimed that "[s]ocial networking has the potential to create an intelligent order in the current chaos by letting you manage how public you make yourself and why and who can contact you" (Black, "The Perils and Promise of Online Schmoozing." *BusinessWeek* Online, February 20, 2004). Similarly, in announcing "more privacy options" and settings that users could control, Facebook's official blog stated: "Today, we are introducing privacy changes that work towards our goal of giving you the control you need in order to share information comfortably on Facebook" ("More Privacy Options," available at http://blog.facebook.com/blog.php?post=11519877130. Retrieved on August 27, 2010). Our results, however suggest that affording more control to users may not necessarily help them to better protect their privacy. The policy implication is that more control may cause people to fall into an "illusion of control" and induce them to reveal more sensitive information.

The peer-reviewed conference version of this paper was accepted in: The Ninth Workshop on the *Economics of Information Security* (WEIS), June 2010, Harvard University, http://weis2010.econinfosec.org/.

The Journal version has been submitted to: *Information Systems Research*.

Authors

Laura Brandimarte is a Ph.D. candidate at Carnegie Mellon University in the Heinz College's Ph.D. program in Public Policy and Management. Her research focuses on the economics of privacy and on the use of behavioral economics to study privacy-related decision making. She has presented papers at conferences including the Ninth Workshop on the Economics of Information Security, and the INFORMS 2009 Annual Meeting.

Before starting her Ph.D. program, Brandimarte lived in her home town of Rome where she completed her undergraduate studies at the University of Rome "La Sapienza." She later received a Master of Science in Economics at the London School of Economics. Currently an instructor of the distance course in Economic Analysis at Heinz College, she previously served as an intern at the European Investment Bank in Luxembourg, and as a consultant for the Italian Federation of Cooperative Banks.

Alessandro Acquisti and **George Loewenstein** biographies listed on page 6 and 7.

Standardizing Privacy Notices
An Online Study of the Nutrition Label Approach

Patrick Gage Kelley, Lucian Cesca, Joanna Bresee and Lorrie Faith Cranor

Full paper available at: www.futureofprivacy.org/the-privacy-papers/

Executive Summary

Through an iterative design process, including focus groups and a laboratory study, we developed a standardized, tabular "nutrition label" for online privacy policies. We tested this standardized format, two variants, and two real-world policy formats in a large, online user study to demonstrate that this label helps consumers.

Introduction & Motivation

Website privacy policies are intended to assist consumers. By notifying them of what information will be collected, how it will be used, and with whom it will be shared, consumers are, in theory, able to make informed decisions. These policies are also meant to inform consumers of the choices they have in managing their information: whether use of their information or sharing with third parties can be limited, and if it is possible to request modification or removal of their information.

However, policies are commonly long, textual explanations of data practices, usually written by lawyers to protect companies against legal action. It has been established through numerous studies that people do not read privacy policies [PLI01] and make mistaken assumptions based upon seeing that a site has a link to a privacy policy [Turow05]. A recent study estimated that if consumers were somehow convinced to read the policies of all the companies they interact with, it would cost an estimated $365 billion per year in lost productivity [McDonald08]. In addition, research has shown that consumers do not actually believe they have choices when it comes to their privacy. Based solely on expectations, they believe there are no options for limiting or controlling companies' use of their personal information [KCG06].

In short, today's online privacy policies are failing consumers because (1) finding information in them is difficult; (2) consumers do not understand that there are differences between privacy policies; and (3) policies take too long to read.

Methodology

We surveyed the literature around standardized and legislated labels [Belser94, Byrd-Bredbenner01, Drichoutis06, FDA94], including nutritional, pharmaceutical, energy ratings and water conservation, to form a baseline from which to design our privacy label. Through four different focus groups and several informal design investigations we refined our prototypes.

To follow up our design process and testing with a full-scale experiment, we tested five privacy policy formats: standardized table, standardized short table (both standardized tables are shown above in Figure 1), standardized short text (Figure 2), full policy text and layered text. Three of these formats are standardized and were created by our lab. Of these, two are tabular and one is textual. Two explicitly describe absent information and one presents it in the context of the policy. Each of these formats is followed by the same list of definitions of privacy terms.

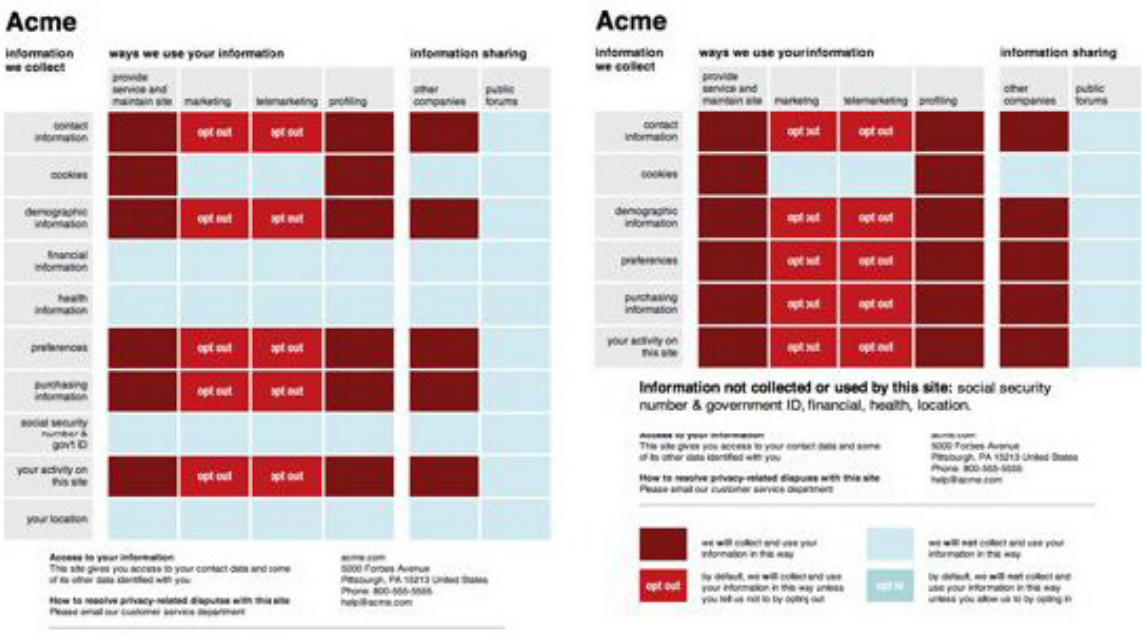

Figure 1. An example of a standardized table is shown on the left and a standardized short table on the right. While both formats contain the legend (bottom right), it is displayed only on the right here due to space constraints.

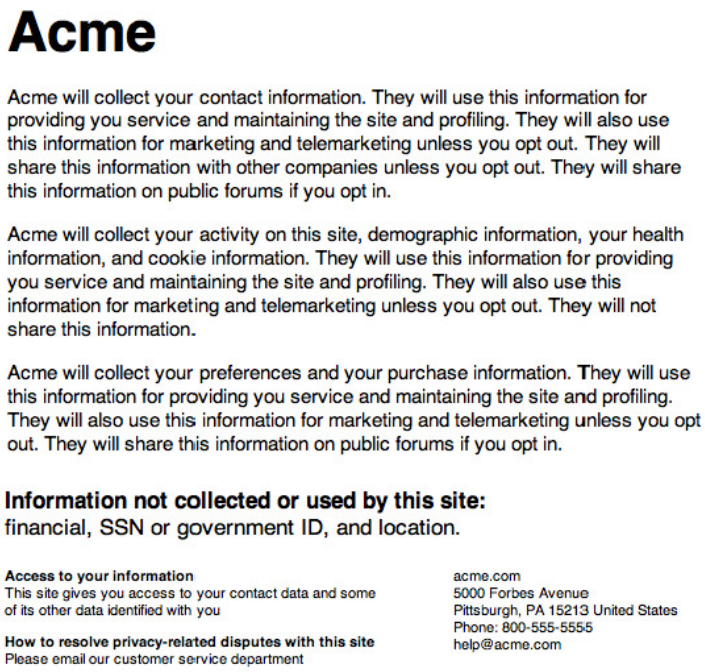

Acme

Acme will collect your contact information. They will use this information for providing you service and maintaining the site and profiling. They will also use this information for marketing and telemarketing unless you opt out. They will share this information with other companies unless you opt out. They will share this information on public forums if you opt in.

Acme will collect your activity on this site, demographic information, your health information, and cookie information. They will use this information for providing you service and maintaining the site and profiling. They will also use this information for marketing and telemarketing unless you opt out. They will not share this information.

Acme will collect your preferences and your purchase information. They will use this information for providing you service and maintaining the site and profiling. They will also use this information for marketing and telemarketing unless you opt out. They will share this information on public forums if you opt in.

Information not collected or used by this site:
financial, SSN or government ID, and location.

Access to your information
This site gives you access to your contact data and some of its other data identified with you

How to resolve privacy-related disputes with this site
Please email our customer service department

acme.com
5000 Forbes Avenue
Pittsburgh, PA 15213 United States
Phone: 800-555-5555
help@acme.com

Figure 2. An example of the standardized short-text format.

We conducted an online user study using Amazon's Mechanical Turk, which offers workers the ability to perform short tasks and get compensated. We paid $0.75 for the successful completion of our approximately 15-minute study. We developed a custom survey-management tool called Surveyor's Point to facilitate our data collection.

We conducted our study and completed the analysis with 764 participants, randomly assigned to the five formats described above. We chose a between-subjects design, meaning participants saw only one of the five formats we tested. Participants in each condition followed the same protocol; only the format differed.

The study included seven groups of questions. It began with basic demographic and Internet usage questions. Next, users were asked to perform simple information-finding tasks in which the question could be answered by looking at a specific row or column. Users were then asked to perform complex information-finding tasks, where the question requires an interaction between a row and a column, or two paragraphs of text. Users were then asked questions about the enjoyability and difficulty of conducting these tasks on a 7-point Likert scale. Next, users were shown two privacy policies in the same format and asked to perform information-finding tasks that required comparing the two policies. Finally users were asked to rate the enjoyability and difficulty of the comparison tasks.

Results

We scored each participant on a scale from 0 to 15, based on how many of the 15 questions they answered correctly, and averaged those scores across conditions. We present these aggregate results in Figure 3. This summary shows a large divide between the standardized and non-standardized formats. The three standardized formats are shown in light blue; while the two real-world text policies are shown in dark red. The standardized policies significantly outperformed the full-text policy in response accuracy.

The standardized formats also significantly outperformed the full policy text in overall time, and the layered format was also significantly faster than the full text policy. The standardized formats, on average were 26 to 32% faster than the full text policy, and 22% faster than the layered text policy.

For the single-policy tasks, participants across the board reported that they felt "confident in my understanding of what I read of Acme's privacy policy." However, for the comparison tasks, participants showed a preference for the standardized formats over the full policy text.

The comments provided by participants at the end of the study provide insights into their enjoyment. Participants who saw the full policy text described privacy policies as "torture to read and understand" and likened them to "Japanese Stereo Instructions." On the other hand, participants in the standardized-format conditions were more complimentary.

For example, one participant wrote: "This layout for privacy policies is MUCH more consumer-friendly. I hope this becomes the industry standard."

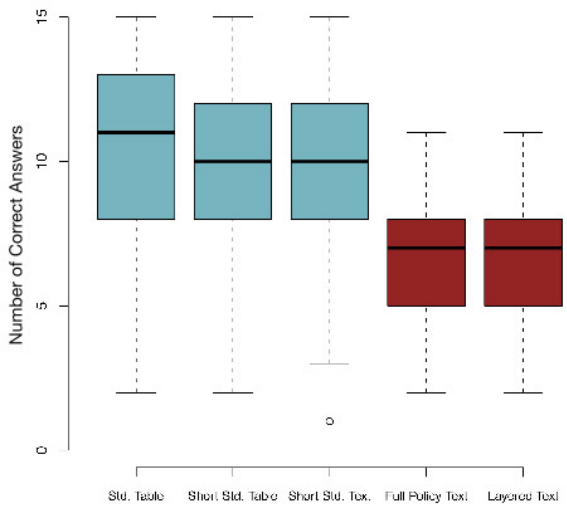

Figure 3: Overall accuracy results for each of the five policy formats.

Conclusions

The final label design allows for information to be found in the same place every time. It removes wiggle room and complicated terminology by using four standard symbols that can be compared easily. It allows for quick, high-level visual feedback by looking at the overall intensity of the page; can be printed; can fit in a browser window; and has a glossary of useful terms attached. People who have used it to find privacy information rated it as pleasurable. They not only rated it better than the natural language policy, but actually rated it enjoyable to use.

The three standardized formats that were designed with usability in mind performed significantly better across a variety of measures than the full-text and layered-text policies that currently exist online today. The large amount of text in full-text policies and the necessity to drill down through a layered policy to the full policy to understand specific practices lengthens the amount of time and effort required to understand a policy. Additionally, more complex questions about data practices frequently require reading multiple sections of these text policies and understanding the way different clauses interact, which is not an easy task. We have shown here that it is not solely table-based formats, but holistic standardization that leads to success. Our standard policies left no room for erroneous, wavering or unclear text, serving as a concise textual alternative to tabular formats.

There is still room for further study and improvement. Complex information-finding tasks and policy- comparison tasks proved difficult. Future work should continue to concentrate on not just how to present policy information, but also on how to facilitate comparisons.

Acknowledgments

The design team was led by Patrick Gage Kelley and included Joanna Bresee, Aleecia McDonald, Robert Reeder, Sungjoon Steve Won and Lorrie Cranor. Thanks to Lucian Cesca, Cristian Bravo-Lillo, Robert McGuire, Daniel Rhim, Norman Sadeh, Clare-Marie Karat and Janice Tsai.

This work was supported in part by U.S. Army Research Office contract DAAD19-02-1-0389 ("Perpetually Available and Secure Information Systems") to Carnegie Mellon University's CyLab, by NSF Cyber Trust grant CNS-0627513, by Microsoft through the Carnegie Mellon Center for Computational Thinking, FCT through the CMU/Portugal Information and Communication Technologies Institute, and the IBM OCR project on Privacy and Security Policy Management.

References

Belser, B. Designing the Food Label: Nutrition Facts. AIGA Journal. 1994.

Byrd-Bredbenner, C., Alfieri, L., Wong, A., and Cottee, P. The Inherent Educational Qualities of Nutrition Labels. Family and Consumer Sciences Research Journal, Vol 29, No 3, March 2001 265-280.

Drichoutis AC, Lazaridis P, Nayga RM.Consumers' use of nutritional labels: a review of research studies and issues. Acad Marketing Sci. Rev, no. 9.2006.

Kleimann Communication Group, Inc. Evolution of a Prototype Financial Privacy Notice. 2006. http://www.ftc.gov/privacy/privacyinitiatives/ftcfi nalreport060228.pdf.

McDonald, A, and Cranor, L.I/S: A Journal of Law and Policy for the Information Society 2008 Privacy Year in Review issue.

Privacy Leadership Initiative. Privacy Notices Research Final Results, November 2001, http://www.ftc.gov/bcp/workshops/glb/supporting/harris%20results.pdf.

Turow, J. Feldman, L., and Meltzer, K. Open to Exploitation: American Shoppers Online and Offline. The Annenberg Public Policy Center. 2005. http://www.annenbergpublicpolicycenter.org/NewsDetails.aspx?myId=31.

U.S. Food and Drug Administration. "Guide to Nutrition Labeling and Education Act Requirements" 1994. http://www.fda.gov/ICECI/Inspections/InspectionGuides/ucm074948.htm

Authors

Patrick Gage Kelley is a Computation, Organizations and Society Ph.D. student in the CyLab Usable Privacy and Security (CUPS) Lab at Carnegie Mellon University. He graduated from the Rochester Institute of Technology with a B.S. in Computer Engineering and a B.A. in Creative Writing and English Literature. While at RIT, he worked for one year with the National Security Agency at the Laboratory for Physical Sciences. At Carnegie Mellon Kelley has worked on projects including Anti-Phishing Phil, a game that teaches users to identify and avoid phishing attacks; Locaccino, a privacy-centered location-sharing system; and the User-Controllable Policy Learning project, which includes a machine learning system for assisting users with privacy setting changes.

Lucian Cesca graduated from Carnegie Mellon in 2010 with a B.S. in Computational and Applied Mathematics. He is currently working for Epic Systems, a healthcare IT firm based near Madison, WI.

Joanna Bresee graduated from Carnegie Mellon in December 2009 with an M.S. in Human Computer Interaction. She is currently a Senior Research Associate at NASA Ames Research Center.

Lorrie Faith Cranor is an Associate Professor of Computer Science and of Engineering and Public Policy at Carnegie Mellon University, where she also serves as director of the CyLab Usable Privacy and Security Laboratory. She has authored more than 80 research papers on online privacy, phishing and semantic attacks, spam, electronic voting, anonymous publishing and usable access control. She has played a key role in building the usable privacy and security research community, having co-edited Security and Usability and founded the Symposium On Usable Privacy and Security. Cranor has served on a number of boards including the Electronic Frontier Foundation Board of Directors. Technology Review magazine named her one of the top 100 innovators 35 or younger. Previously a researcher at AT&T-Labs Research, she also taught at the Stern School of Business at New York University.

How Different are Young Adults from Older Adults When it Comes to Information Privacy Attitudes and Policies?

Chris Hoofnagle, Jennifer King, Su Li and Joseph Turow

Full paper available at: www.futureofprivacy.org/the-privacy-papers/

Executive Summary

Media reports teem with stories of young people posting salacious photos online, writing about alcohol-fueled misdeeds on social networking sites, and publicizing other ill-considered escapades that may haunt them in the future. These anecdotes are interpreted as representing a generational shift in attitude toward information privacy. Many commentators therefore claim that young people "are less concerned with maintaining privacy than older people are." Surprisingly, though, few empirical investigations have explored the privacy attitudes of young adults. This report is among the first quantitative studies to evaluate young adults' attitudes. It demonstrates that the picture is more nuanced than portrayed in the popular media.

In this telephonic (wireline and wireless) survey of Internet using Americans (N=1000), we found that large percentages of young adults (18-24year-olds) agree with older Americans in regard to concerns about online privacy, norms and policy suggestions. In several cases, there are no statistically significant differences between young adults and older age categories. Where there were differences, more than half of the young adult respondents did answer in a similar direction as older adults. There clearly is social significance in that large numbers of young adults agree with older Americans on issues of information privacy.

For example, a large majority of young adults:

- Have refused to give information to a business in cases where they felt it was too personal or not necessary;
- Believe anyone who uploads a photo of themselves to the Internet should get their permission first, even if it was taken in public;
- Believe there should be a law that gives people the right to know all the information websites know about them; and
- Believe there should be a law that requires websites to delete all stored information about an individual.

In view of these findings, why would so many young adults act in social networks and elsewhere online in ways that would seem to offer quite private information to all comers? A number of answers present themselves, including suggestions that people 24 years and younger approach cost-benefit analyses related to risk differently than do individuals older than 24.

An important part of the picture, though, must surely be our finding that higher proportions of 18-24 year-olds believe incorrectly that the law protects their privacy online and offline more than it actually does. Forty-two percent of young Americans answered all of our five online privacy questions incorrectly;

88% answered only two or fewer correctly. The problem is even more pronounced when presented with offline privacy issues – post hoc analysis showed that young Americans were more likely to answer no questions correctly than any other age group.

We conclude that young adult Americans hunger for increased privacy, even while they participate in an online reality that is optimized to increase revelation of personal data.

This study is the second in a series of three examining the gaps between privacy law and consumers' perceptions of protections. In our first study, "Americans Reject Tailored Advertising and Three Activities that Enable It," we found that contrary to what many marketers claim, most adult Americans (66%) do not want marketers to tailor advertisements to their interests. Moreover, when Americans are informed of three common ways that marketers gather data about people in order to tailor ads, even higher percentages-- between 73 and 86 %-- say they would not want such advertising.

The third study, to be released fall 2010, will focus on the differences in privacy attitudes among social network users and non-users.

Authors

Chris Jay Hoofnagle is director of the Berkeley Center for Law & Technology's Information Privacy Programs, and a senior fellow to the Samuelson Law, Technology & Public Policy Clinic. He is an expert in information privacy law.

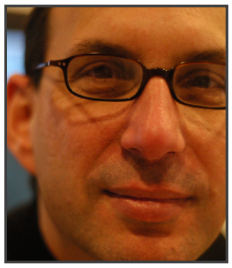

Joseph Turow is Robert Lewis Shayon Professor of Communication at the University of Pennsylvania's Annenberg School for Communication. He has authored eight books, edited five and written more than 100 articles on mass media industries. He is a Fellow of the International CommunicationAssociation and was named a Distinguished Scholar by the National Communication Association. He was recently awarded an Astor Visiting Lectureship by Oxford University. His books include *Niche Envy: Marketing Discrimination in the Digital Age*, *Breaking Up America: Advertisers and the New Media World, The Hyperlinked Society: Questioning Connections in the Digital Age and The Wired Homestead*.

Jennifer King, a Ph.D. candidate at the UC Berkeley School of Information, has worked as a researcher at the Samuelson Law, Technology, and Public Policy Clinic at UC Berkeley's School of Law. Her research areas include information privacy and security, privacy in online social networks, ubiquitous computing, and sensor networks (e.g. video surveillance, RFID), and usability and privacy. King has co-authored three reports about Californians' privacy attitudes, available at SSRN.com.

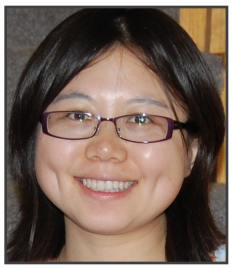

Su Li is the Statistician of Empirical Legal Studies at University of California, Berkeley School of Law. She received her Ph.D. in Sociology and a Master's degree in Mathematical Methods for Social Sciences from Northwestern University. An expert in quantitative methodology, Li served as an Assistant Professor of Sociology at Wichita State University before joining Berkeley Law. Her research interests include law and society, gender and social inequality, economic sociology, social network analysis and the sociology of education.

Privacy and Regulatory Innovation: Moving Beyond Voluntary Codes

Ira Rubinstein

Full paper available at: www.futureofprivacy.org/the-privacy-papers/

Executive Summary

Privacy policy in the U.S. has long relied on a combination of sectoral law, market forces and self-regulation. Over the years, the Department of Commerce and the Federal Trade Commission (FTC)—with the strong backing of industry—have expressly favored a self-regulatory approach. They argue that it can protect privacy in a more flexible and cost-effective manner than direct regulation without impeding the rapid pace of innovation in Internet-related businesses. On the other hand, critics of self-regulation emphasize its shortcomings, including weak or incomplete realization of Fair Information Practice Principles (FIPPs), inadequate incentives to ensure wide scale industry participation, ineffective compliance and enforcement mechanisms, and an overall lack of transparency. Typically, these same critics call upon Congress to intervene in the online marketplace by enacting comprehensive privacy legislation. Under this enforcement model of regulation, Congress would define substantive privacy requirements for commercial firms based on FIPPs and grant the FTC both regulatory and enforcement powers.

The opposing sides in the privacy debate tend to treat self-regulation and direct government regulation as if they were mutually exclusive options from which policy makers have to choose: either one or the other. But self-regulation is a highly malleable term that may encompass a wider variety of policy instruments. Thus, it is better to think of self-regulation and direct regulation as opposing ends of a regulatory continuum, with most regulatory schemes falling somewhere in the middle. Rather than attack or defend self-regulation, this Article explores a "co-regulatory" approach in which industry enjoys considerable scope in shaping self-regulatory guidelines, while government retains general oversight authority to approve and enforce these guidelines. In particular, this hybrid approach builds on the idea of a *privacy safe harbor.*

In general, a safe harbor is a regulatory strategy under which a federal statute recognizes differences in industry performance explicitly by treating qualifying firms more favorably than non-qualifying firms. In other words, safe harbors shield or reward firms if they engage in desirable behavior as defined by statute. Favorable treatment for better performing firms might include immunity from liability, protection from certain penalties, exemptions from certain requirements, and/or permission to engage in certain desired behaviors. The key point is that eligibility for the benefits conferred by a safe harbor are contingent upon a participating firm meeting a higher standard of performance than what is otherwise required of firms covered by the relevant statute.

Congress first created a privacy safe harbor in the Children's Online Privacy Protection Act of 1998 (COPPA). The COPPA safe harbor establishes an alternative means of compliance for operators that follow self-regulatory guidelines issued by an industry representative and approved by the FTC, subject to a notice and comment procedure, and it seeks to facilitate industry self-regulation in two ways: first, by granting enforcement-related benefits (operators that comply with approved self-regulatory

guidelines are deemed to be in compliance with the law); and second, by allowing greater flexibility in the development of self-regulatory guidelines in a manner that takes into account industry-specific concerns and technological developments. FTC approval of a COPPA safe harbor program turns on whether self-regulatory guidelines (1) meet or exceed statutory requirements; (2) include an effective, mandatory mechanism for the independent assessment of compliance with the guidelines; and (3) contain effective incentives to ensure compliance with the guidelines. In practice, the COPPA safe harbor programs have met with success mainly in terms of complementing the FTC's own enforcement efforts. However, the COPPA safe harbor also suffers from two serious shortcomings: first, a very low rate of participation (because deemed compliance is too weak an incentive to persuade many firms to bear the costs of joining a safe harbor program); and, second, a lack of regulatory flexibility (all of the approved self-regulatory programs have nearly identical requirements to those of the COPPA statute).

This Article proposes that in considering new privacy legislation, Congress should redesign the COPPA safe harbor program in several critical respects using a more co-regulatory approach. The basic idea underlying this proposal is that a co-regulatory approach could more effectively use both sticks and carrots as incentives. In other regulatory settings such as environmental, sticks typically include a threat of stricter regulations or imposition of higher pollution fees, whereas carrots might take the form of more flexible regulations, recognition of better performance by the government, and cost-savings such as exemptions from mandatory reporting, or easier and quicker permitting. Firms that demonstrate high performance avoid these sticks and/or enjoy these carrots. What does this mean in the privacy setting, and why might it attract higher rates of industry support than did the COPPA safe harbor program?

For many years, advocacy groups and privacy scholars have favored a private right of action and liquidated damages as enforcement mechanisms in any new privacy legislation. Not surprisingly, industry has argued that such remedies are both unnecessary and ineffective. This suggests that an excellent stick might be devised around a tiered liability system. Under this new approach, privacy legislation would allow civil actions and liquidated damages awards against firms that engage in prohibited practices and did not participate in an approved safe harbor program. In sharp contrast, compliance with approved self-regulatory guidelines would not only serve as a safe harbor in any enforcement action but exempt program participants from civil lawsuits and monetary penalties. Other sticks for non-participating firms might include broader opt-in requirements; external and independent audits of regulatorycompliance and mandatory reporting to the FTC; and much stricter requirements for firms engaged in online behavioral advertising, such as a total ban on the use of sensitive information in behavioral targeting and a data retention limit of one month.

In addition to these sticks, privacy legislation might also offer safe harbor participants a number of carrots including exemptions from civil actions and liquidated damages; cost-savings such as compliance reviews based on self-assessments rather than external audits by an independent third party; government recognition of better performing firms (e.g., an FTC "seal of approval" under which firms that meet safe harbor requirements are duly recognized); government procurement preferences for the products or services of participating firms; and regulatory flexibility in the form of tailored requirements addressed to specific business models such as online behavioral advertising.

In summarizing this new approach to privacy safe harbors, it is important to emphasize that safe harbor benefits should be limited to firms demonstrating superior performance and not be available to other covered entities that merely satisfy default statutory requirements. In other words, a safe harbor provides incentives, in the form of sticks and carrots, but only to firms that meet higher performance standards. As described more fully in the Article, these standards might include data governance practices (i.e., a system for assigning rights and accountabilities within a company for all information-related processes); advanced privacy methodologies (such as the use of development guidelines for building privacy protection into any product or service that uses personal data, a process often referred to as "Privacy by Design"); and industry-wide best practices (such as mandatory privacy training for all staff with privacy responsibilities, providing online guidance on privacy and security issues to employees and consumers, and implementing a complaint-handling procedure).

This Article develops the ideas described above in considerable detail and has three parts. Part I begins by analyzing the rise and fall of self-regulation as the FTC's preferred approach to online privacy in the five-year period ending in 2000, when it finally recommended that Congress enact a basic level of online privacy protection. It then examines the Commission's shift in 2001, under Chairman Tim Muris, to an enforcement-based agenda designed to remedy specific harms, as well as the FTC's renewed interest in self-regulation as the best way to handle the privacy concerns raised by behavioral advertising. Finally, it considers the arguments of privacy scholars and economists for and against self-regulation, but finds this debate inconclusive since both sides voice compelling objections without advancing a solution that resolves their differences.

Part II shifts from a more general and abstract discussion of self-regulation to three case studies involving co-regulatory solutions. The first is a "weakly mandated" industry effort aimed at online behavioral advertising practices; the second is a more "strongly mandated" safe harbor program resulting from joint governmental efforts to ensure data flows between Europe and the U.S.; and the third is a statutorily mandated safe harbor under COPPA.

Part III begins by assessing these case studies against five criteria—completeness, free rider problems, oversight and enforcement, transparency, and formation of regulatory norms—and concludes that despite several weaknesses, statutory safe harbors such as COPPA offer a superior form of self-regulation. Next, it describes a more collaborative, flexible and performance-based approach to self-regulation, drawing on critical insights from environmental regulation. It discusses one specific set of policy tools known as environmental covenants and applies this learning to the use of privacy covenants (which would enable the FTC to experiment with policy innovations such as Privacy by Design) and to a revamped version of statutory safe harbors (which might result in much broader industry participation).

This re-designed version of privacy safe harbors would ensure both the superior performance of participating firms, and a baseline level of compliance by all firms, thereby allowing the FTC to devote its scarce enforcement resources to the most egregious or systemic privacy abuses. Therefore, the Article strongly recommends that Congress adopt these new tools to better protect online consumer privacy.

Author

Ira Rubinstein is an Adjunct Professor of Law at New York University School of Law and a Senior Fellow at the Information Law Institute. His research interests include Internet privacy, electronic surveillance law, online identity, Internet security and software liability. He lectures and publishes widely on issues of privacy and security, and has testified before Congress on several occasions, including last July on new privacy legislation before the U.S. House of Representatives, Subcommittee on Commerce, Trade, and Consumer Protection. His most recent publications include "Privacy and Regulatory Innovation: Moving Beyond Voluntary Codes," I/S: A Journal of Law and Policy for the Information Society (forthcoming Winter 2011) and "Data Mining and Internet Profiling: Emerging Regulatory and Technological Approaches," co-authored with Ron Lee and Paul Schwartz, 75 U. Chi. L. Rev. 261 (2008).

Previously he spent 17 years in Microsoft's Legal and Corporate Affairs department, most recently as Associate General Counsel in charge of the Regulatory Affairs and Public Policy group, and before that was in private practice in Seattle, specializing in immigration law. Rubinstein has served on the President's Export Council, Subcommittee on Encryption, and on the Editorial Board of the IEEE Security and Privacy Magazine. In 2010, he joined the Board of Directors of the Center for Democracy and Technology. He graduated from Yale Law School in 1985.

Privacy Papers of Notable Mention

To View The Following Papers Visit: www.futureofprivacy.org/the-privacy-papers/

"A Model Regime of Privacy Protection"
By Daniel Solove and Chris Hoofnagle

"The Deidentification Dilemma: A Legislative and Contractual Proposal"
By Robert Gellman

"Sponsoring Trust in Tomorrow's Technology: Towards a Global Digital Infrastructure Policy"
By John Miller and David Hoffman

"Making Sense of Privacy and Publicity"
By Danah Boyd

"Preemption and Privacy"
By Paul Schwartz

"Privacy by Design 7 Foundational Principles: Implementation and Mapping of Fair Information Practices"
By Ann Cavoukian

"The Eavesdropping Employer: A Twenty-first Century Framework For Employee Monitoring"
By Corey Ciocchetti

"The Boundaries of Privacy Harm"
By M. Ryan Calo

"Americans Reject Tailored Advertising and Three Activities that Enable It"
By Joseph Turow

This report was designed with the environment and cost-effectiveness in mind. It is printed on recovered fiber paper that has no ozone layer threatening emissions and generates no detectable amounts of sulfur, chlorine, nitrogen, or dioxide gases when properly incinerated.

The Future of Privacy Forum (FPF) is a Washington, D.C.-based think tank that seeks to advance responsible data practices. The forum is led by Internet privacy experts Jules Polonetsky and Christopher Wolf, and includes an advisory board comprised of leading figures from industry, academia, law and advocacy groups. FPF was launched in November 2008, and is supported by Adobe, AOL, AT&T, The Better Advertising Project, Bering Media, BlueKai, Comcast, Datran Media, Deloitte, eBay, Facebook, Intel, Lockheed Martin, Microsoft, The Nielsen Company, Qualcomm, TRUSTe, Verizon and Yahoo.

To learn more about FPF, please visit www.futureofprivacy.org