

NECESSARY BUT NOT SUFFICIENT: STANDARDIZED MECHANISMS FOR PRIVACY NOTICE AND CHOICE

LORRIE FAITH CRANOR*

I.	NOTICE AND CHOICE.....	277
II.	PLATFORM FOR PRIVACY PREFERENCES	279
III.	A PRIVACY TAXONOMY.....	282
IV.	PRIVACY NUTRITION LABELS AND PRIVACY ICONS.....	286
V.	ADOPTION AND ENFORCEMENT	295
VI.	OPTING OUT OF ONLINE BEHAVIORAL ADVERTISING	299
VII.	CONCLUSIONS	304

For several decades, “notice and choice” have been key principles of information privacy protection.¹ Conceptions of privacy that involve the notion of individual control require a mechanism for individuals to understand where and under what conditions their personal information may flow and to exercise control over that flow. Thus, the various sets of fair information practice principles and the privacy laws based on these principles include requirements for providing notice about data practices and allowing individuals to exercise control over those practices. Privacy policies and opt-out mechanisms have become the predominant tools of notice and choice. However, a consensus has emerged that privacy policies are poor mechanisms for communicating

* Associate Professor, Computer Science and Engineering & Public Policy and Director, CyLab Usable Privacy and Security Laboratory at Carnegie Mellon University. lorrie@cmu.edu. I presented a very preliminary draft of this paper at Berkeley Law as part of the 4th Annual Privacy Lecture in February 2011. Thanks to respondents Thomas Fetzer and Jennifer Gove, as well as Paul Schwartz and Deirdre Mulligan for their feedback and suggestions. I presented a later draft at the Silicon Flatirons “The Economics of Privacy” Symposium in December 2011. Discussions with conference participants, several former P3P working group members, and members of the Carnegie Mellon CyLab Usable Privacy and Security Laboratory further informed this paper.

1. See Memorandum from Paul M. Schwartz & Daniel Solove on Notice and Choice: Implications for Digital Marketing to Youth prepared for the Second NPLAN/BMSG Meeting on Digital Media and Marketing to Children (June 29-30, 2009), http://digitalads.org/documents/Schwartz_Solove_Notice_Choice_NPLAN_BMSG_memo.pdf

with individuals about privacy.² These policies are long, complicated, full of jargon, and change frequently. If an individual were to read the privacy policy at every website she visited even once per year, she would spend, on average, an estimated 244 hours per year reading privacy policies.³

To make matters worse, visiting a single website today typically involves interactions with multiple parties unknown to the end user. Each of these “third parties” – including advertising networks, analytics providers,⁴ and service providers that assist the website in customizing content for its visitors – collects bits of data about site visitors. Those data bits might be immediately deleted after being used to select a targeted ad, or they might be combined with hundreds of other bits of information about a particular user and stored indefinitely as part of a digital dossier. In this environment, it is nearly impossible for website visitors to determine where their data flows, let alone exert any control over it. Privacy policies for the first-party websites that users interact with are difficult enough for users to understand, but when third-party sites enter the mix, the notion of effective privacy notice becomes completely untenable.

With growing recognition that website privacy policies are failing consumers, numerous suggestions⁵ are emerging for technical mechanisms that would provide privacy notices in machine-readable form, allowing web browsers, mobile devices, and other tools to act on them automatically and distill them into simple icons for end users. Other proposals are focused on allowing users to signal to websites, through their web browsers, that they do not wish to be tracked.⁶ These proposals may at first seem like fresh ideas that allow us to move beyond impenetrable privacy policies as the primary mechanisms of notice and choice. Facilitating transparency and control through easily recognizable symbols and privacy controls that need be set only once are laudable goals. However, in many ways, the conversations around these new proposals are reminiscent of those that took place in the 1990s that led to

2. Fred H. Cate, *The Limits of Notice and Choice*, 8 IEEE SEC. & PRIVACY 59, 59–62 (2010).

3. Aleecia M. McDonald & Lorrie Faith Cranor, *The Cost of Reading Privacy Policies*, 4 ISJLP 560 (2008).

4. For a brief overview of analytics and privacy issues see Paul M. Schwartz, *Privacy, Ethics, and Analytics*, 9 IEEE SEC. & PRIVACY 66, 66–69 (2011).

5. Here are just a few recent examples: TRUSTe has stated an intention to support efforts to develop XML privacy policies (<http://www.truste.com/blog/2010/09/14/more-on-the-problem-with-p3p/>). Mozilla has launched a privacy icons project (https://wiki.mozilla.org/Drumbeat/Challenges/Privacy_Icons). The Interactive Advertising Bureau (IAB) CLEAR Ad Notice project plans to integrate XML privacy notices (<http://www.iab.net/clear>).

6. DO NOT TRACK: UNIVERSAL WEB TRACKING OPT OUT, <http://donottrack.us> (last visited May 1, 2012).

the development of the Platform for Privacy Preferences (“P3P”) standard⁷ and several privacy seal programs.⁸ I was reminded of this when I looked back at my own contribution to a 1997 US Department of Commerce Report, “Privacy and Self-Regulation in the Information Age.”⁹ I outlined several approaches to simplifying notice and choice, including privacy icons, a machine-readable label system, and a system that allowed web browsers to communicate user privacy preferences to websites automatically.

Reviewing other essays from the early days of US online privacy self-regulation reveals more similarities to the present. Indeed the privacy regulatory landscape in early 2012 – in which there had been several recent Congressional hearings on privacy¹⁰ and the privacy community was awaiting reports from both the Federal Trade Commission¹¹ (“FTC”) and the Department of Commerce¹² – bears a

7. For a more complete history of P3P see chapter 4 of LORRIE FAITH CRANOR, WEB PRIVACY WITH P3P 44–61 (2002). For another account of the history and a discussion of related policy issues see Harry Hochheiser, *The Platform for Privacy Preferences as a Social Protocol: An Examination within the U.S. Policy Context*, 2 ACM TRANSACTIONS ON INTERNET TECH. 276-306 (Nov. 2002). For a more recent account see also Ari Schwartz, Ctr. for Democracy & Tech., Looking Back at P3P: Lessons for the Future (2009), http://www.cdt.org/files/pdfs/P3P_Retro_Final_0.pdf.

8. One of the early proposals for the eTRUST privacy seal program (later changed to TRUSTe) involved the seal provider offering three different levels of trust marks to describe three different types of data sharing practices. Each mark would have its own icon. Esther Dyson describes this in Esther Dyson, *Labels and Disclosure Part II: Privacy*, RELEASE 1.0: ESTHER DYSON’S MONTHLY REPORT, Feb. 19, 1997, at 1, 6, available at <http://cdn.oreilly.com/radar/r1/02-97.pdf>.

9. Lorrie Faith Cranor, *The Role of Technology in Self-Regulatory Privacy Regimes*, in PRIVACY AND SELF REGULATION IN THE INFORMATION AGE, NAT’L TELECOMMS. & INFRASTRUCTURE ADMIN., U.S. DEP’T OF COMMERCE, 185, 185–191 (1997), available at <http://www.ntia.doc.gov/reports/privacy/selfreg5.htm#5B>.

10. 2011 Congressional privacy hearings included Senate Judiciary Committee Subcommittee on Privacy, Technology and the Law hearings on mobile privacy and health privacy; a full Senate Judiciary Committee hearing on the Electronic Communications Privacy Act; a full Senate Commerce Committee hearing on privacy and data security; and House Energy and Commerce Committee hearings on Understanding Consumer Attitudes About Privacy, which focused on consumer attitudes about privacy, protecting children’s privacy, Internet privacy, and the effectiveness of privacy controls.

11. FED. TRADE COMM’N, PROTECTING CONSUMER PRIVACY IN AN ERA OF RAPID CHANGE: A PROPOSED FRAMEWORK FOR BUSINESSES AND POLICYMAKERS PRELIMINARY FTC STAFF REPORT (2010), available at <http://www.ftc.gov/os/2010/12/101201privacyreport.pdf>. For the final report, see FED. TRADE COMM’N, PROTECTING CONSUMER PRIVACY IN AN ERA OF RAPID CHANGE: RECOMMENDATIONS FOR BUSINESSES AND POLICYMAKERS, FTC REPORT (2012), available at <http://ftc.gov/os/2012/03/120326privacyreport.pdf>.

12. INTERNET POLICY TASK FORCE, U.S. DEP’T OF COMMERCE, COMMERCIAL DATA PRIVACY AND INNOVATION IN THE INTERNET ECONOMY: A DYNAMIC POLICY FRAMEWORK (2010), available at http://www.ntia.doc.gov/files/ntia/publications/iptf_privacy_greenpaper_12162010.pdf. For the final report, see THE WHITE HOUSE, CONSUMER DATA PRIVACY IN A NETWORKED WORLD: A FRAMEWORK FOR PROTECTING PRIVACY AND PROMOTING INNOVATION IN THE GLOBAL DIGITAL ECONOMY (2012), available at

sharp resemblance to Esther Dyson's description of the situation in February 1997:

Currently, the government is indeed paying substantial attention to privacy issues on several fronts. The Federal Trade Commission is conducting a long-term Privacy Initiative and is planning a privacy workshop to study technical tools and self-regulatory models to protect privacy At the same time, the Commerce Department . . . is compiling a report on the issues around privacy self-regulation. "As a general matter, " says NTIA chief counsel Barbara Wellbery, "we favor self-regulation, but self-regulation with teeth How do you handle enforcement? What role can technology play in all of this?" . . . there are also several bills pending in Congress.¹³

Dyson goes on to describe the emerging effort to develop P3 (as P3P was called at the time) and eTRUST (as TRUSTe was called at the time). Those efforts proceeded, but in 2012 we find ourselves in more or less the same place we were in 1996 when these efforts were launched. The Federal Trade Commission is once again calling on companies to "increase the transparency of their data practices"¹⁴ and the Department of Commerce is calling "for multi-stakeholder efforts to produce voluntary, enforceable codes of conduct."¹⁵

In this paper I first review the idea behind notice and choice and user empowerment as privacy protection mechanisms. Next I review lessons from the development and deployment of P3P as well as other efforts to empower users to protect their privacy. I begin with a brief introduction to P3P, and then discuss the privacy taxonomy associated with P3P. Next I discuss the notion of privacy nutrition labels and privacy icons and describe our demonstration of how P3P policies can be used to generate privacy nutrition labels automatically. I also discuss studies that examined the impact of salient privacy information on user behavior. Next I look at the problem of P3P policy adoption and enforcement. Then I discuss problems with recent self-regulatory programs and privacy tools in the online behavioral advertising space. Finally, I argue that while standardized notice mechanisms may be necessary to move beyond impenetrable privacy policies, to date they have failed users and they will continue to fail users unless they are accompanied by usable mechanisms for exercising meaningful choice and appropriate means of enforcement.

<http://www.whitehouse.gov/sites/default/files/privacy-final.pdf>.

13. See Dyson, *supra* note 8, at 14.

14. FED. TRADE COMM'N, *supra* note 11, at 69.

15. U.S. DEP'T OF COMMERCE, *supra* note 12, at vii.

I. NOTICE AND CHOICE

In his often-cited 1967 book, *Privacy and Freedom*, Alan Westin defines privacy as “the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others.” While a range of privacy definitions exists, this definition, which focuses on individual control, is the definition around which most modern data privacy principles and laws are based. Westin explains that “each individual is continually engaged in a personal adjustment process in which he balances the desire for privacy with the desire for disclosure and communication of himself to others, in light of the environmental conditions and social norms set by the society in which he lives.”¹⁶

This notion of privacy as control is reflected in the 1973 U.S. Department of Housing, Education, and Welfare Fair Information Practices (“FIPs”). The FIPs require notice of data collection and use and provide the right to control the use of data for purposes beyond which it was collected.¹⁷ The 1980 OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data contain similar principles, drawing out explicitly the principles of collection and use limitation.¹⁸ In Europe, privacy laws are based closely on the OECD Guidelines. In the United States, we have a patchwork of privacy laws, some of which incorporate FIPs, but no comprehensive data protection laws. Instead we have relied mostly on a self-regulatory, market-driven approach to privacy protection, based loosely on the FIPs. However, U.S. industry has simplified the FIPs considerably and distilled them into the concept of “notice and choice,” which is often interpreted to mean: “As long as a company provides notice of its privacy practices, and people have some kind of choice about whether to provide the data or not, then privacy is sufficiently protected.”¹⁹

The US Federal Trade Commission held a series of privacy workshops beginning in June 1996, which led to a report to Congress in 2000. In this report, the FTC endorsed the industry’s simplified view of the FIPs, which eliminates the collection and use limitation principles.²⁰ Thus, posting a privacy policy and allowing individuals to opt-out of

16. ALAN F. WESTIN, *PRIVACY AND FREEDOM* 7 (1967).

17. COMM. ON AUTOMATED PERSONAL DATA SYS. U.S. DEP’T OF HEALTH, EDUC, AND WELFARE, RECORDS, COMPUTERS, AND THE RIGHTS OF CITIZENS 41–42 (1973).

18. DIRECTORATE FOR SCI., TECH. & INDUS., OECD, OECD GUIDELINES ON THE PROTECTION OF PRIVACY AND TRANSBORDER FLOWS OF PERSONAL DATA (1980), http://www.oecd.org/document/18/0,3746,en_2649_34223_1815186_1_1_1_1,00.html.

19. Schwartz & Solove, *supra* note 1.

20. U.S. FED. TRADE COMM’N, *PRIVACY ONLINE: FAIR INFORMATION PRACTICES IN THE ELECTRONIC MARKETPLACE*, REPORT TO CONGRESS (2000), *available at* <http://www.ftc.gov/reports/privacy2000/privacy2000.pdf>.

some uses of their data passes for providing notice and choice and complying with the FIPs. For most of the next decade, the FTC encouraged a market-based approach to privacy, but acknowledged that more effort would be required to make this approach successful. Howard Beales, then Director of the FTC Bureau of Consumer Protection, said in a 2002 speech, “First, privacy notices should be viewed as a means of facilitating competition over privacy practices. Their goal should be to help consumers understand what information is collected about them and what is done with that information, not to simply scare consumers into opting out of information sharing.” But Beales also emphasized that the need for privacy notices to be understandable by consumers, and suggested taking an approach similar to the approach taken to develop nutrition labels, which involved extensive consumer testing and public education. At the same time he warned against “rigidly prescribed disclosure formats” and “adding additional notices and forms to those consumers are already receiving.”²¹

By 2010, the FTC staff was even more vocal about the shortcomings of privacy notices and proclaimed that “the notice-and-choice model, as implemented, has led to long, incomprehensible privacy policies that consumers typically do not read, let alone understand.” FTC staff recommended that “Privacy notices should be clearer, shorter, and more standardized, to enable better comprehension and comparison of privacy practices.” Beyond the shortcomings of privacy notices, the FTC staff went on to question the market-driven, notice and choice approach to privacy: “industry efforts to address privacy through self-regulation have been too slow, and up to now have failed to provide adequate and meaningful protection.”²²

The notice and choice model relies on the presence of the elements necessary for a meaningful decision making process. Recall Westin’s description of the “personal adjustment process” that individuals engage in continuously.²³ Westin envisions an individual actively involved in a decision-making process. To exert control, this individual must be aware of the consequences of both disclosing and not disclosing personal information, and must have the ability to effectively govern whether information is disclosed and how it will be used. In practice, individuals often lack complete information about the consequences of information disclosure as well as mechanisms for ensuring that their information is

21. Howard Beales, Dir., Bureau of Consumer Prot., Remarks on the Privacy Notices and the Fed. Trade Comm’n’s 2002 Privacy Agenda (Jan. 24, 2002), (transcript available at <http://www.ftc.gov/speeches/other/privacynotices.shtm>).

22. Press Release, Fed. Trade Comm’n, FTC Staff Issues Privacy Report, Offers Framework for Consumers, Businesses, and Policymakers (Dec. 1, 2010).

23. WESTIN, *supra* note 16.

disclosed or used only in the ways they desire.²⁴ Even when information is available, processing this information may be more burdensome than is feasible for a continual process that is supposed to occur in the background, as a secondary task as we go about our daily living. While we may easily close our blinds and lower our voices to adjust our personal privacy as needed, making online privacy decisions can be a much more difficult and time-consuming process. Therefore, for the past 15 years, notice and choice proponents have advocated the use of “user empowerment” tools to help provide meaningful and accessible notice and choice.²⁵

II. PLATFORM FOR PRIVACY PREFERENCES

P3P is a user empowerment tool that was developed in response to Congressional and FTC interest in online privacy in the mid-1990s and concerns that it was unrealistic to expect consumers to read long online privacy policies at every website they visit. Instead, early proponents of P3P described web browsers that could read privacy policies, negotiate with websites, and take actions on their users’ behalf without interfering with the web browsing experience.²⁶ P3P was envisioned as a tool that could facilitate a market for privacy, enabling individuals to shop around for websites that would match their privacy preferences, refusing to do business with those they found unacceptable, and perhaps accepting payments or discounts in exchange for data.²⁷

After nearly two years of informal discussions, in 1997 the World Wide Web Consortium (“W3C”) launched a five-year process that led to the publication of the P3P 1.0 specification in 2002.²⁸ The original idea for P3P involved a protocol in which web browsers would negotiate with websites over privacy on behalf of their users. The negotiation protocol was not included in the final specification, largely due to the added implementation complexity and lack of interest from industry, but also due to concerns that negotiations would not benefit consumers.²⁹ A

24. LORRIE FAITH CRANOR, *Privacy Policies and Privacy Preferences*, in SECURITY AND USABILITY: DESIGNING SECURE SYSTEMS THAT PEOPLE CAN USE 448 (Lorrie Faith Cranor & Simson Garfinkel, eds., 2005).

25. CRANOR, *supra* note 9, at 185-191.

26. DYSON, *supra* note 8.

27. *Id.* at 2. See also LAWRENCE LESSIG, CODE AND OTHER LAWS OF CYBERSPACE 160 (Basic Books, 1999); William McGeeveran, Note, *Programmed Privacy Promises: P3P and Web Privacy Law*, 76 N.Y.U. L. Rev. 1812, 22-23 (2001).

28. Massimo Marchiori, ed. *The Platform for Privacy Preferences 1.0 (P3P1.0) Specification, W3C Recommendation*, W3C (16 April 2002), <http://www.w3.org/TR/P3P/>.

29. Cranor and Resnick provide a theoretical analysis of P3P negotiation that show that it is likely to encourage websites to adopt strategies that will not benefit consumer privacy. Lorrie Faith Cranor & Paul Resnick, *Protocols for Automated Negotiations with Buyer Anonymity and Seller Reputations*, 2 NETNOMICS 1, 1-23 (2000).

subsequent P3P 1.1 effort produced a working draft that included additional P3P vocabulary elements, backwards-compatible syntax changes, and plain-language definitions of P3P elements.³⁰ However, the P3P working group was closed in 2006 due to lack of industry participation, and P3P 1.1 was never finalized.

P3P 1.0 provides an XML format for website privacy policies, and a protocol for locating and retrieving these policies and associating them with online resources. The XML format encodes the P3P “vocabulary,” a privacy taxonomy that was the subject of much debate and disagreement during the P3P development process. The P3P protocol is fairly simple, designed so that no special software would be required for web servers to comply with P3P. Websites can become P3P-enabled simply by placing P3P files at designated locations on their servers. Most of the complexity associated with the P3P protocol centers around performance optimizations designed to reduce the number of P3P requests that user agents must make to locate and fetch up-to-date P3P policies.

The P3P 1.0 specification also describes a P3P “compact policy” format for providing a summary of the privacy policy for cookies that can be transferred in an HTTP header. The compact policy was intended as a supplement to a full P3P policy, designed to allow browsers to evaluate quickly the policies associated with cookies. The P3P specification requires sites using compact policies to provide accompanying full P3P policies.³¹

P3P user agent tools have been integrated into the Microsoft Internet Explorer 6, 7, 8, and 9 web browsers,³² as well as Netscape 7.³³ P3P was never implemented for Firefox,³⁴ Safari, or Chrome, although a number of prototype plug-ins and extensions have been developed.³⁵ In addition, a variety of P3P authoring tools³⁶ have been developed as well

30. Rigo Wenning & Matthias Schunter, eds. *The Platform for Privacy Preferences 1.1 (P3P1.1) Specification*, W3C Working Group Note, W3C (13 November 2006), <http://www.w3.org/TR/P3P11/>.

31. MARCHIORI, *supra* note 28, at § 4.

32. Microsoft, *Privacy in Microsoft Internet Explorer 6*, [http://msdn.microsoft.com/en-us/library/ms537343\(v=vs.85\).aspx](http://msdn.microsoft.com/en-us/library/ms537343(v=vs.85).aspx) (last visited June 2, 2012).

33. Harish Dhurvasula, Dave Barrowman, and Steve Morse, *Technical Issues in Implementing P3P in Netscape 7.0*, W3C Workshop on the Future of P3P (2002), <http://www.w3.org/2002/p3p-ws/pp/netscape.html>.

34. Some Mozilla developers were planning to implement P3P support at one point, and had begun writing the code. However, this project appears to have been abandoned. See Tom Lendacky, *The Platform for Privacy Preferences (P3P)* (2002), <http://www-archive.mozilla.org/projects/p3p/>.

35. See for example, Privacy Bird for Chrome, a P3P user agent available for free download from the Chrome Web Store.

36. One of the most popular P3P authoring tools is the P3P Policy Editor distributed for free by IBM. See alphaWorks Community, <http://www.alphaworks.ibm.com/tech/p3peditor> (last visited May 2, 2012).

as prototype P3P user agents.³⁷

The Internet Explorer P3P implementation is probably the most widely used P3P tool given the widespread use of IE. However, it appears that most IE users are completely unaware of P3P.³⁸ P3P functionality is associated with three user interface components in IE, although the interface does not explicitly mention P3P in any of those places. First, in the View menu, users have the option of viewing a “Privacy Report.” Clicking on this option causes IE to check whether a website has a full P3P policy. If IE finds a P3P policy, it fetches it and translates the XML code into English (or the appropriate language for that version of IE), using the technical definitions of P3P elements found in the P3P 1.0 specification. Second, the default cookie setting in IE (the medium setting) bases third-party cookie-blocking decisions on P3P compact policies. Third-party cookies without compact policies are blocked. IE analyzes any compact policies it finds associated with cookies and determines whether or not the policies are “satisfactory.” Those third-party cookies found to have unsatisfactory compact policies are blocked.³⁹ Finally, when IE blocks cookies it places a small icon in the bottom right area of the browser chrome that looks like a do-not-enter sign overlapping an eye. Most users do not seem to have any idea what the icon means. However, those who click on the icon are shown a list of blocked cookies. Users can click on a link for each blocked cookie to display a privacy report if there is a full P3P policy associated with it.

Microsoft’s decision to base third-party cookie-blocking decisions in Internet Explorer on P3P compact policies led to widespread adoption of P3P among advertising networks and other companies making substantial use of third-party cookies. P3P was adopted by about a third of the most popular websites, but never saw widespread adoption beyond popular sites and those that use third-party cookies.⁴⁰

From the beginning, a number of privacy advocates opposed P3P on the grounds that industry groups were using it “as an excuse to delay the progress of genuine enforceable privacy rights in the US.”⁴¹ P3P

37. I have been involved in the development of an IE browser helper object called Privacy Bird, <http://privacybird.org>, and a P3P-enabled search engine called Privacy Finder, <http://privacyfinder.org>. Privacy Finder demonstrates the use of P3P to help users select privacy-protective sites from among search results. It also integrates a privacy “nutrition label” generated automatically from P3P policies. *Privacy Nutrition Labels*, CYLAB USABLE PRIVACY AND SEC. LAB., <http://cups.cs.cmu.edu/privacyLabel/>.

38. While I know of no formal studies, my informal polls of hundreds of audience members at talks I have given suggests that outside of groups of privacy experts, almost nobody has heard of P3P or has any idea what the IE blocked-cookie icon represents.

39. Microsoft, *supra* note 32.

40. Lorrie Faith Cranor, Serge Egelman, Steve Sheng, Aleecia M. McDonald, & Abdur Chowdhury, *P3P Deployment on Websites*, 7 Elec. Commerce Research and Applications 3, 274-93 (Autumn 2008), available at <http://lorrie.cranor.org/pubs/p3p-deployment.html>.

41. Jason Catlett, *An Open Letter to P3P Developers*, JUNKBUSTERS.COM (Sep. 13,

supporters responded that P3P was complementary to other regulatory and self-regulatory privacy efforts and was not intended as a substitute for enforceable privacy rights.⁴²

III. A PRIVACY TAXONOMY

The P3P vocabulary provides a taxonomy of privacy practices. P3P policies begin with some general assertions about the location of human-readable policies and opt-out mechanisms, and website contact information. Websites must choose between six disclosures about the type of access they will provide to their users' identified data. In addition, they may optionally describe one or more mechanisms for resolving privacy-related disputes. For each mechanism described, sites provide a URL and description, and classify it as one of four dispute-resolution mechanism types and optionally associate it with any of the three defined remedy types. The rest of the P3P policy consists of one or more "statements" that describe a set of data and the practices that apply to that data set. Each statement includes data categories (from a list of 17 possible data categories), purposes (from a list of 12 possible purposes), recipients (from a list of six possible types of recipients), and retention (from a list of five possible types of retention policies). Statements can also include a human-readable description and enumerate specific data elements. In addition, attributes can be used to indicate that certain purposes are done on an opt-in or opt-out basis, or that certain types of data are optional. P3P also includes an extension mechanism that can be used to add additional vocabulary components.

The P3P vocabulary has long been criticized simultaneously for being too complicated, and for not being expressive enough for companies to accurately represent their privacy practices.⁴³ Indeed, there is a tension between the need to develop a standard that is simple enough to be practically implemented and expressive enough to capture nuances of privacy practices. This tension is exacerbated by the fact that end users and companies often have different ideas about what details of privacy practices are important to represent. Some companies have even criticized P3P for exposing the "gory detail" of their privacy practices.⁴⁴

1999), <http://www.cfp2000.org/papers/catlett.pdf>; See also *Pretty Poor Privacy: An Assessment of P3P and Internet Privacy*, ELEC. PRIVACY INFORMATION CENTER AND JUNKBUSTERS (June 2000), <http://epic.org/reports/prettypoorprivacy.html>.

42. OFFICE OF THE INFO. AND PRIVACY COMM'R, THE CTR. FOR DEMOCRACY & TECH., ONTARIO, P3P AND PRIVACY: AN UPDATE FOR THE PRIVACY COMMUNITY, <http://www.ipc.on.ca/images/Resources/p3p.pdf> (2000).

43. Hochheiser, *supra* note 7.

44. E-Mail from Kenneth Lee and Gabriel Speyer to the Citibank Advanced Dev. Group, *White Paper: Platform for Privacy Preferences Project (P3P) and Citibank*, available at http://www.w3.org/P3P/Lee_Speyer.html.

During the P3P development process, some people argued for a vocabulary with fewer elements, while others argued for a more expressive vocabulary. For example, consider the problem of describing data sharing using the P3P recipient element. Some people argued that all consumers want to know is whether or not a website shares their data. However, corporate representatives argued that companies share data for many reasons and consumers should be given an opportunity to better understand the type of sharing (with the implication that they were likely to find some types of sharing less objectionable than others). Along the way, the working group considered proposals to distinguish between sharing with parent companies, sharing with subsidiaries, sharing with business partners, sharing with companies with similar privacy practices, sharing with companies with unknown privacy practices, and many other combinations. Indeed, at one point I counted about three-dozen different recipient elements that had been proposed. In the end a compromise was reached and six types of recipients were included in the specification.⁴⁵ In my subsequent work on P3P user agents, I have not found utility in exposing more than three types of recipients to end users (shares, hosts a public forum for users to share, doesn't share except with agents).⁴⁶

In principle, a more expressive taxonomy might be preferred because it captures more information. If users don't care about this level of detail, a user agent designer can build interfaces that suppress some of the detail.⁴⁷ However, if the detail is available in the taxonomy and in the computer-readable policy, advanced interfaces can expose it to those expert users who are interested in it. Furthermore, over time if a particular practice becomes more important to disclose (perhaps due to new technology or changing regulations), if the detail is available user agents can be updated without requiring the underlying computer-readable policies to change.

So why not develop an extremely detailed taxonomy and rely on user agent designers to distill this detailed information into something more readily understood by users? Despite having been simplified somewhat, the P3P 1.0 vocabulary represents a fairly detailed privacy taxonomy and is thus a good case study to investigate this question. My experience with P3P leads me to identify three problems with this approach, which I describe below. That said, I believe there are tradeoffs

45. CRANOR *supra* note 7, at 198-199.

46. Patrick Gage Kelley, Lucian Cesca, Joanna Bresee, & Lorrie Faith Cranor, *Standardizing Privacy Notices: An Online Study of the Nutrition Label Approach*, CARNEGIE MELLON CYLAB, (2009), available at http://www.cylab.cmu.edu/research/techreports/tr_cylab09014.html.

47. Lorrie Faith Cranor & Joseph Reagle Jr., *Designing a Social Protocol: Lessons Learned from the Platform for Privacy Preferences Project*, TELECOMM. POLICY RESEARCH CONFERENCE (Sept. 27-29, 1997), <http://www.w3.org/TR/NOTE-TPRC-970930/>.

and a detailed taxonomy may still have merit despite these challenges.

First, we have found that P3P policy authors have trouble distinguishing between some of the P3P vocabulary elements and there are some elements that are commonly used incorrectly. For example, the “historical” purpose is designed for government agencies to explain that they keep data for historical purposes required by law. However, many P3P policy authors have misunderstood this element and used it on commercial websites where it is clearly inappropriate.⁴⁸ In addition, the P3P vocabulary includes four different profiling elements, plus a fifth “tailoring” element used to indicate that a site is customized for a user during a particular session without building a profile of that particular user. The distinctions between the four profiling elements allow sites to differentiate between identified and anonymous profiling, and between profiling used for analysis purposes versus decision-making purposes (a distinction that may yet prove important as people try to define what tracking should be prohibited in a “do-not-track” system). However, these distinctions seem to be difficult for policy authors to understand, even if they are sometimes useful.⁴⁹

Second, the complicated P3P vocabulary, including most importantly the syntax rules about how the various elements can be combined, has added complexity to user agent implementations. Indeed, the P3P implementation in Internet Explorer suffers from implementation bugs, perhaps due in part to this complexity. For example, IE reports the data categories from each statement in the privacy report, but omits any data elements mentioned explicitly by name. This omission may mislead users when websites declare individual data elements rather than categories.⁵⁰ When my students have developed prototype P3P user agent implementations as class projects I have observed them making similar types of errors. This is not an excuse for software errors and, compared to other protocols, P3P is not really that complicated a protocol to implement. Nonetheless it is worth keeping in mind that complexity adds overhead to software development.

Third, because each user agent developer may make different decisions about how to simplify the P3P vocabulary, P3P policy authors have to test their policies on all P3P user agents in order to see how they will look to end users. The fact that some P3P user agents do not provide faithful representations of some P3P policy elements, as noted above, makes this all the more problematic. The current dearth of widely used P3P user agents means this particular problem is not really much of an

48. CRANOR, *supra* note 40.

49. CRANOR, *supra* note 7, at 94-95.

50. Lorrie F. Cranor & Joel R. Reidenberg, *Can User Agents Accurately Represent Privacy Notices?*, 30 TELECOMM. POLICY RESEARCH CONFERENCE (Sept. 28-30, 2002) <http://intel.si.umich.edu/tprc/archive-search-abstract.cfm?PaperID=65>.

issue in practice, but it was a concern that has been raised repeatedly.⁵¹ To encourage a more standardized approach among P3P user agents a section on user agent guidelines, including plain language translation of all the P3P vocabulary elements was added to P3P 1.1.⁵²

While the detailed P3P vocabulary has posed challenges, the much simpler P3P compact policy syntax has also been problematic. P3P compact policies simplified P3P in a way that reduces expressiveness such that companies have had difficulty expressing their policies accurately without making them appear to use data much more extensively than they actually do. Compact policies are particularly problematic for companies that must rely on them to avoid IE cookie blocking. Recall that a full P3P policy allows sites to declare multiple P3P “statement” elements. Sites tend to use elements to group together data types that are used in a common way. For example, a site might have one statement for the types of data it collects in its server logs and another statement for the types of data it collects from users who register on the site. Compact policies do not have a way to form statement groups. Therefore, all data types and purposes get thrown together. Thus a site has no way to explain in its compact policy that it will share anonymous preference information but not personally identified contact information; the compact policy will state that the site collects preference information and contact information and shares all of it.⁵³

The P3P vocabulary has stood the test of time somewhat, providing most of the elements needed to express 2012 data practices even though it was developed more than ten years earlier. One of the elements that wasn’t useful when P3P was developed but has become useful now is the location category, used to indicate that a website collects a user’s precise location information (through GPS or other location-tracking technology). The “location” category is part of the P3P vocabulary because some members of the working group anticipated that it would become common for websites to request precise location information. On the other hand, the P3P working group did not anticipate the extent of peer-to-peer personal information sharing that is done through social networks. The “public” recipient allows websites to express the fact that peer-to-peer information sharing takes place. However, P3P does not allow any finer level of detail such as whether users can restrict sharing to only their friends or to friends of friends. It is not clear whether this is

51. Daniel M. Schutzer, *Citigroup P3P Position Paper*, W3C (Sept. 20, 2002), <http://www.w3.org/2002/p3p-ws/pp/citigroup.html>.

52. Wenning & Schunter, *supra* note 30, at §6.0 User Agent Guidelines.

53. In 2006 a minor change to the compact policy syntax was proposed in P3P 1.1 to address this problem and significantly improve expressivity. However, P3P 1.1 was never finalized and this syntax was not adopted in Internet Explorer. See Wenning & Schunter, *supra* note 30, at §4.2.10.

a level of detail that is needed in P3P or not.

While P3P is frequently criticized for being insufficiently expressive, few critics have pointed to concrete examples of where more expressiveness is needed in practice. During the P3P 1.1 discussions a jurisdiction element was proposed in response to concerns that sites could not comply with the European Union directive without disclosing what jurisdiction they were in. In addition, a primary purpose element was proposed to address the concern that P3P purposes focus on secondary use of data and do not encode the primary purpose associated with each data element. However, the FTC staff has recently suggested that companies streamline their privacy notices and focus on secondary data uses rather than those primary data uses that are either obvious to the user or necessary and commonly accepted practices.⁵⁴

I have also observed a number of problems related to P3P syntax that allows the same practice to be described in multiple ways, ambiguous definitions for some P3P policy elements, and a fairly convoluted syntax. For example, a website may convey that they use cookies by including the “miscdata” element with the state management mechanism category or by including the cookies element. Either way, the site is required to indicate all categories of data that might be linked to a cookie, but may exclude data that is potentially linkable but not actually linked to a cookie. The definition of “linked” was sufficiently confusing in P3P 1.0 that a new section was added to the draft P3P 1.1 Specification to explain the meaning of “linked” and “linkable”.⁵⁵

IV. PRIVACY NUTRITION LABELS AND PRIVACY ICONS

The development of the P3P specification was motivated in part by the desire to allow consumers to make decisions based on website privacy policies without having to read privacy policies at every site they visit. Online privacy policies are notoriously confusing and difficult to understand.⁵⁶ Because of their complex language and inconsistent structure and format, they are also extremely difficult to compare.

54. FTC Preliminary Staff Report 2010, *supra* note 11, at vi (“[I]t is reasonable for companies to engage in certain commonly accepted practices – namely, product and service fulfillment, internal operations such as improving services offered, fraud prevention, legal compliance, and first-party marketing. Some of these practices, such as where a retailer collects a consumer’s address solely to deliver a product the consumer ordered, are obvious from the context of the transaction, and therefore, consent for them is inferred. Others are sufficiently accepted – or necessary for public policy reasons – that companies need not request consent to engage in them. By clarifying those practices for which consumer consent is unnecessary, companies will be able to streamline their communications with consumers, reducing the burden and confusion on consumers and businesses alike.”).

55. See section 1.3.4 Linked and Linkable data of Wenning & Schunter, *supra* note 30.

56. Carlos Jensen & Colin Potts, *Privacy Policies as Decision-Making Tools: An Evaluation of Online Privacy Notices*, 6 SIGCHI 471, 472 (2004).

Besides the P3P project, a number of other efforts have attempted to address this concern through standardized privacy notices or icons.

One of the first efforts to develop a standardized privacy notice was the multilayer privacy notice project organized by the Center for Information Policy Leadership at the Hunton and Williams law firm.⁵⁷ Multilayer privacy notices involved a standardized one-page top layer with links into a full privacy policy. The standardized top layer included standardized section headings and rough guidelines for what should be included in each section. However, it offered companies a lot of flexibility to determine for themselves exactly what should go in each section and the terminology to use. This flexibility was appealing to companies, but made it difficult for consumers to use the top layer effectively. In our user testing we found that participants did not know where to look to find specific pieces of information and could not determine when they needed to click through to the full policy to find information that was omitted from the top layer.⁵⁸

In 2009, seven Federal Agencies jointly announced a model privacy notice for financial organizations that are required by the Gramm-Leach-Bliley Act to send annual privacy notices to their customers.⁵⁹ This notice was developed over a five-year period, with the assistance of consumer research commissioned by the Agencies.⁶⁰ The model privacy notice takes the form of a table.

Outside of the privacy realm, other types of consumer communications are much easier for consumers to understand because they have been standardized and use summary views to provide the most salient information at a glance. For example, nutrition labels on food packages offer a brief standardized format, as well as a complete list of ingredients. In the United States, standardized food nutrition labels were mandated by the Nutrition Labeling and Education Act of 1990 (“NLEA”).⁶¹ Studies of the impact of NLEA have found mostly small effects.⁶² However, they have found that nutrition labels, as well as

57. *Ten Steps to Develop a Multilayered Privacy Notice*, CENTER FOR INFO. POLICY LEADERSHIP (Mar. 15, 2007), http://www.hunton.com/files/tbl_s47Details/FileUpload265/1405/Ten_Steps_whitepaper.pdf.

58. Kelley et al., *supra* note 46.

59. Final Model Privacy Form Under the Gramm-Leach-Bliley Act, 74 Fed. Reg. 62,890-62,994 (Dec. 1, 2009).

60. KLEIMANN COMM’N GRP., INC., *EVOLUTION OF A PROTOTYPE FINANCIAL PRIVACY NOTICE 1* (2006), http://www.ftc.gov/privacy/privacy_initiatives/ftcfinalreport060228.pdf; *see also* ALAN LEVY & MANOJ HASTAK, *CONSUMER COMPREHENSION OF FINANCIAL PRIVACY NOTICES: A REPORT ON THE RESULTS OF THE QUANTITATIVE TESTING 1* (2008), http://www.ftc.gov/privacy/privacy_initiatives/Levy-Hastak-Report.pdf.

61. *Guide to Nutrition Labeling and Education Act Requirements*, FDA (Aug. 1994), <http://www.fda.gov/ICECI/Inspections/InspectionGuides/ucm074948.htm>.

62. Andreas C. Drichoutis, et al., *Consumers’ Use of Nutritional Labels: A Review of Research Studies and Issues*, 9 ACAD. MKTG. SCI. REVIEW 1, 2 (2006),

calorie information on restaurant menus, can be particularly useful to people who are dieting.⁶³ Other standardized consumer communications include the US FDA Drug Facts label⁶⁴ on pharmaceuticals and energy labels on appliances.⁶⁵

Inspired by nutrition labels,⁶⁶ we used an iterative design process to develop and test a privacy nutrition label. Our evaluations suggest that our most recent design (shown in Figure 1) allows consumers to find information more quickly and accurately than traditional text privacy policies. We believe that the nutrition label has a number of advantages over traditional privacy policies. For example, it is shorter and easier to read than a traditional text privacy policy, and its standardized tabular format allows users to learn where to look to find information in a consistent location and facilitates comparison between policies. In addition, the use of colored symbols allows users to get an overview of a policy at a glance from observing the overall color intensity of a policy.⁶⁷

Nutrition label research has found that even when nutrition labels are readily available, consumers still do not read them every time they make a purchase. However, the labels still play an important role for consumers who seek them out so that they can eat according to specific dietary restrictions (whether due to a medical condition, a desire to lose weight, a preference for organic food, or other factors).⁶⁸ The availability of nutrition labels also allows journalists and thought leaders to obtain ready access to nutrition information that they can use to educate the public and policy makers.⁶⁹ Privacy nutrition labels are likely to play a similar role. The FTC staff suggest that “the public posting of privacy notices is especially valuable to consumer and privacy advocacy groups, regulators, and those consumers who want to learn more about a company’s overall privacy practices.” They go on to state that “although privacy policies may not be a good tool for communicating with most consumers, they still could play an important role in promoting

<http://www.amsreview.org/articles/drichoutis09-2006.pdf>.

63. Julie S. Downs et al., *Strategies for Promoting Healthier Food Choices*, 99 AM. ECON. REV. 159, 159-164 (2009).

64. *New OTC Drug Facts Label*, FDA CONSUMER MAGAZINE (July 2002), http://permanent.access.gpo.gov/lps1609/www.fda.gov/fdac/features/2002/402_otc.html.

65. *The Energy Label* (2007), www.energyrating.gov.au.

66. The FTC’s efforts to standardize financial privacy notices were also inspired, in part, by nutrition labels. See BEALES, *supra* note 21.

67. Kelley et al., *supra* note 46.

68. Drichoutis, Lazaridis & Nayga, *supra* note 62, at 4.

69. The *Men’s Health* magazine series “Eat This, Not That!” highlights particularly unhealthy foods and healthier alternatives, and provides highlights from nutrition label information for the foods they review. Other publications often report on the *Men’s Health* findings. For example, a feature on “Worst Drinks in America” made headlines around the world. 20 Worst Drinks in America, MENSHEALTH (2010), <http://eatthis.menshealth.com/slideshow/20-worst-drinks-america-2010> (last visited Apr. 23, 2012).


transparency, accountability, and competition among companies on privacy issues – but only if the policies are clear, concise, and easy-to-read. Thus, companies should improve their privacy policies so that interested parties can compare data practices and choices across companies.”⁷⁰ The privacy nutrition label may be a viable approach.


Bell Group							
information we collect	ways we use your information				information sharing		
	to provide service and maintain site	marketing	telemarketing	profiling	other companies	public forums	
contact information		opt in			opt out		
cookies							
demographic information		opt in			opt out		
financial information							
health information							
preferences							
purchasing information		opt in			opt out		
social security number & gov't ID							
your activity on this site		opt in			opt out		
your 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

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

 we will collect and use your information in this way

 we will not collect and use your information in this way

 **opt out** by default, we will collect and use your information in this way unless you tell us not to by opting out


 **opt in** by default, we will not collect and use your information in this way unless you allow us to by opting in

FIGURE 1: AN EXAMPLE PRIVACY NUTRITION LABEL

While our privacy nutrition label format could be implemented manually, we designed it so that it could be generated automatically by websites with P3P policies. We have incorporated automatically generated privacy nutrition labels into the privacy reports produced by our Privacy Finder search engine (<http://privacyfinder.org>).⁷¹ Ideally, the nutrition label would also be linked directly to opt-out mechanisms and

70. FTC Privacy Report Dec. 2010, *supra* note 22, at vii, 70.

71. For details about how Privacy Finder generates the nutrition label, see The Privacy Label, PRIVACY FINDER, http://www.privacyfinder.org/about_label (last visited Apr. 23, 2012).

facilitate automating the opt-out process. However, as there are no standards for opt-out mechanisms, the best we can currently do is provide a link to the page on each website where the user can get information about opting out.

In some of our early attempts to develop a privacy nutrition label, we tried to capture all P3P elements into a standardized tabular format.⁷² This proved overwhelming to users. We eventually ended up with a simplified table in which we collapsed together similar data categories, purposes, and recipients and did not attempt to display each statement in a policy separately. We reduced the 17 purpose elements to 10 rows in our table, plus an explicit note about sites that use the “other” purpose (which is not represented by a table row because it requires a human-readable explanation). We collapsed the 12 data categories into 4 columns plus explicit notes about the rare historical and other-purposes. In addition, we collapsed the 6 recipient elements into 2 columns and omitted the recipient elements that indicate that information is used by the company that collected it and its agents and by delivery companies, since these almost always apply.

Here is a summary of how we collapsed the rows and columns:⁷³

- We combined physical contact information and online contact information into a single “contact information” row.
- We combined the categories for preferences, political and religious affiliations, and messages you send or post on the site into a single “preferences” row.
- We combined the categories for website login IDs, click stream data, activities on the site, and computer information into a single “your activity on this site” row.
- We also combined four purposes into a column for “provide services and maintain site.”
- We combined the four profiling purposes into a column for “profiling.”
- We combined the three recipients that involve sharing data with other companies into a single “other companies” column.

Certainly some detail has been lost by collapsing all of these rows and columns, and it is not difficult to imagine situations where a user would care deeply about what happens to some types of information that have now been bundled together and not care about others. However, we believe that in most cases the added details are unnecessary to display.

72. Robert W. Reeder, et al., *A User Study of the Expandable Grid Applied to P3P Policy Visualization*, WORKSHOP ON PRIVACY IN THE ELECTRONIC SOCIETY (WPES) (Oct. 28, 2008), available at <http://lorrie.cranor.org/pubs/wpes24reeder.pdf>.

73. The Privacy Label, *supra* note 64.

Future enhancements may include the ability to mouse-over or click on individual table cells to reveal more detailed information, including information about the individual data categories that have been collapsed. Because there is additional information in the underlying P3P policy, it can be pulled out automatically and provided to users who want more detailed information.

Even after collapsing the rows and columns as described above, the nutrition label may still provide more information than many users actually want most of the time. In addition, it is sufficiently large that it needs to be on a page by itself. Thus, there is also a need for small privacy icons that could be integrated into web pages or in web browsers to allow users to get a quick understanding of a privacy policy without having to click through to the privacy nutrition label. In our user studies of Privacy Finder, we found that a 5-point privacy meter (shown in Figure 2), represented by a string of green and white boxes, helped users quickly find the websites with the best privacy policies and influenced user's decisions about where to make purchases.⁷⁴ A small set of icons highlighting some of the practices users may be most concerned about or a thumbnail image of the nutrition label itself might also work. Because there are so many dimensions to privacy, the challenge here is coming up with something that is fairly simple and focuses on the most relevant dimensions.

74. Serge Egelman et al., *Timing Is Everything? The Effects of Timing and Placement of Online Privacy Indicators*, CHI 2009 (Apr. 2009) (Proceedings of the SIGCHI Conference on Human Factors in Computing Systems), available at <http://www.guanotronic.com/~serge/papers/chi09a.pdf>.

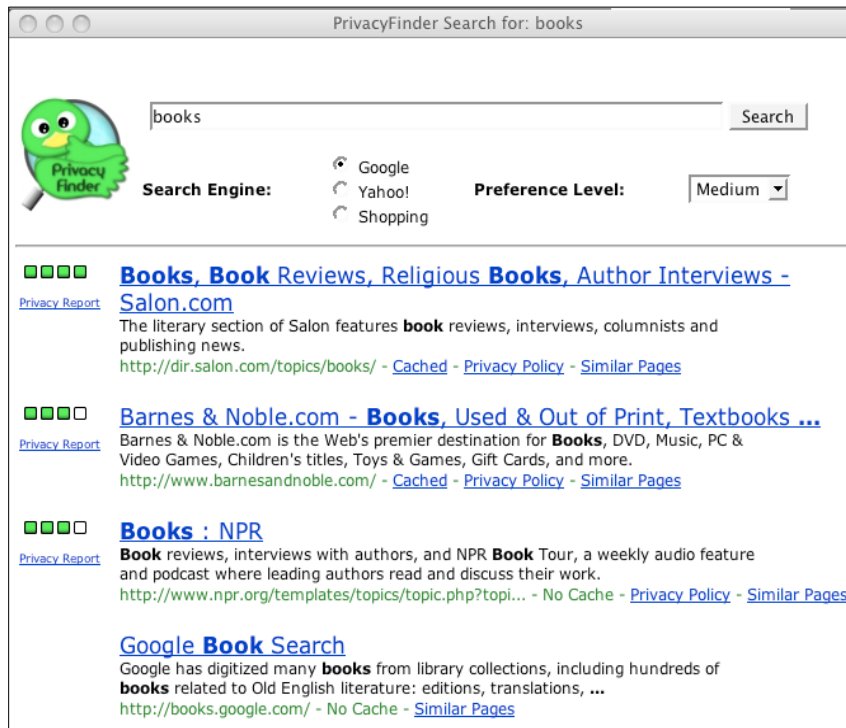


FIGURE 2: PRIVACY FINDER SEARCH RESULTS ANNOTATED WITH GREEN PRIVACY METERS

We conducted a series of laboratory studies to investigate the effect of our privacy meter icons in search results. For example, we invited 89 participants to our lab to participate in an evaluation of a new online shopping search engine. Our study included four conditions in which timing and placement of the privacy icons varied across conditions. Each study participant was asked to search for two items using our search engine and select a vendor from which to purchase each item using their own credit card and providing their personal billing information. We selected one item that we expected to raise privacy concerns, and one item that was not expected to raise particular privacy concerns beyond the concerns associated with providing contact and billing information to a website. We carefully controlled the search results and the prices of the items at the websites that appeared in the first page of search results. This required making arrangements with these websites to make minor adjustments to their prices for the items we asked our participants to buy. We selected merchants such that those with the best privacy policies were the most expensive. This allowed us to measure the impact of privacy icons on purchasing behavior. We found that in the condition without privacy icons, most participants made their purchases from the least expensive websites. However, in the conditions where privacy indicators were present, a significant number of participants paid extra to

buy the items at the more privacy-protective web sites. This effect was especially pronounced for the privacy-sensitive item. In addition, the privacy icons were most influential when presented to users in the search results, before they visited the merchant websites.⁷⁵

In 2009, three University of California, Berkeley graduate students developed a set of privacy icons as part of an academic project.⁷⁶ They used these black-and-white circular icons to rate company privacy policies on a website they created called knowprivacy.org. The KnowPrivacy icons, shown in Figure 3, include five icons for types of data collected, five icons for general data practices, and three icons for data sharing. While these icons do not correspond exactly to the rows and columns in the privacy nutrition label, there is a lot of similarity. Travis Pinnick, one of the students who developed the KnowPrivacy icons, later joined the staff of TRUSTe and developed a privacy short notice design that incorporated variations on the KnowPrivacy icons with red and green rings to indicate privacy invasive and protective practices.⁷⁷ After doing some user testing of this design he concluded, “Icon Design is not as important as category selection and taxonomic presentation. Several users commented that initially the purpose of the short notice is to educate, and as long as the icons made reasonable sense in the context of the categories they would eventually come to be associated with their intended meanings.”⁷⁸

75. Egelman et al., *supra* note 74.

76. JOSHUA GOMEZ, TRAVIS PINNICK & ASHKAN SOLTANI, KNOWPRIVACY FINAL REPORT (2009), available at http://www.knowprivacy.org/report/KnowPrivacy_Final_Report.pdf.

77. Travis Pinnick, *Privacy Short Notice Design*, TRUSTE BLOG (Feb. 17, 2011), <http://www.truste.com/blog/2011/02/17/privacy-short-notice-designpart-i-background/>.

78. Travis Pinnick, *Layered Policy Design*, TRUSTE BLOG (May 20, 2011), <http://www.truste.com/blog/2011/05/20/layered-policy-and-short-notice-design/>.














TYPE OF DATA COLLECTED	GENERAL DATA PRACTICES	DATA SHARING
 contact: name, mailing address, email, or phone number	 ad customization: user data may be used for the purpose of customizing advertising	 affiliates: affiliates and subsidiaries bound by the same privacy practices
 computer: IP address, browser type, or operating system	 third party tracking: site allows third parties to place advertisements that may track user behavior	 contractors: third party contractors bound by the same privacy practices
 interactive: browsing behavior or search history	 public display: service allows users to contribute information which may be displayed publicly	 third parties: third parties not subject to same data practices
 financial: account status or activity, credit information, or purchase history	 user control: users allowed to access and correct personal data collected	
 content: contents of personal communications, stored documents or media	 data retention: explicitly stated duration of retention for personal data collected	

FIGURE 3: KNOW PRIVACY ICONS⁷⁹

In 2010, Aza Raskin led an effort at Mozilla to develop a set of privacy icons. To date there has been only an “alpha release” of these icons and Mozilla has not announced plans to use these icons in their web browser.⁸⁰ The icon set includes 10 circular icons (as well as a number of variants to indicate how long information is retained), shown in Figure 4. Half of the icons have green borders to represent good privacy, and half have red borders to represent bad privacy. Three of the green icons are divided vertically by a dotted line that is intended to represent that a type of data sharing does not take place. The icons include some concepts absent from P3P and the other icon sets, including “data is given to law enforcement only when legal process is followed.”

79. *Policy Coding Methodology*, KNOWPRIVACY, http://knowprivacy.org/policies_methodology.html (last visited Apr. 23, 2012).

80. Aza Raskin, *Privacy Icons: Alpha Release*, AZA ON DESIGN (Dec. 27, 2010), <http://www.azarask.in/blog/post/privacy-icons/>.




FIGURE 4: AZA RASKIN'S PRIVACY ICONS, ALPHA RELEASE⁸¹

V. ADOPTION AND ENFORCEMENT

Arguably, the largest barrier to P3P adoption has not been problems with the P3P vocabulary or difficulties with the technical mechanisms, but rather lack of incentives to adopt. As Dyson observed in 1997, "Industry disclosure schemes often founder without strong government/public pressure. Otherwise, companies are simply too busy to adopt them, and customers don't factor the information disclosed into their buying habits."⁸² By the time the P3P specification was released in 2002, government pressure had subsided and industry had largely lost interest in P3P.

Besides a set of companies who adopted P3P because they were positioning themselves as privacy leaders, most of the adopters decided to implement P3P in order to prevent IE6 from blocking their cookies. However, over time we began to observe that many of these companies did not appear to be making serious efforts to implement P3P, and instead were offering minimal policies designed to prevent IE cookie blocking. Initial signs of this were the large number of sites that

81. *Id.*

82. Dyson, *supra* note 8, at 15.

implemented P3P compact policies without the corresponding full policies, and an unexpectedly large number of sites that had syntax errors in their P3P policies.⁸³

More recently, our research has found that a large fraction of sites adopting P3P compact policies have misrepresented their privacy practices, most likely in an effort to prevent IE from blocking their cookies.⁸⁴ We collected compact policies from 33,139 websites and used automated techniques to detect syntax errors and inconsistencies in 11,176 of them, including 134 TRUSTe-certified websites and 21 of the top 100 most-visited sites. Upon further investigation, we discovered thousands of sites that had identical erroneous policies and traced these policies to a Microsoft support website⁸⁵ and several blog posts that recommended posting these policies to prevent cookie blocking.

We also found sites that posted compact policies that bore little resemblance to proper compact policy syntax and were clearly meant to circumvent IE. For example, Amazon posted a compact policy containing the single made-up token “AMZN” and Facebook posted a compact policy containing only tokens for the disputes and remedies elements (no data categories, purpose, recipients, access, or retention tokens, which are required for a valid policy). During a preliminary study in 2009 we observed that the Facebook compact policy contained the single made-up token “HONK.”

We also discovered that when IE analyzes compact policies to determine whether they are satisfactory, it simply looks for combinations of tokens that appear on a list of unsatisfactory tokens. IE apparently does not test the compact policy to determine whether it is syntactically valid. As a result, compact policies that consist entirely of made-up tokens will never be flagged as unsatisfactory.

In March 2011, a class action lawsuit was filed that alleged, among other things, that Amazon’s P3P compact policy circumvents web browser privacy settings so that cookies are not blocked.⁸⁶ The case was dismissed in December 2011, largely because the plaintiffs did not allege

83. Egelman et al., *supra* note 40.

84. Pedro Giovanni et al., *Token Attempt: The Misrepresentation of Website Privacy Policies through the Misuse of P3P Compact Policy Tokens* 4, WORKSHOP ON PRIVACY IN THE ELECTRONIC SOCIETY (WPES) (Oct. 2010), available at http://www.cylab.cmu.edu/files/pdfs/tech_reports/CMUCyLab10014.pdf.

85. Session variables are lost if you use FRAMESET in Internet Explorer 6, MICROSOFT SUPPORT (Apr. 2006), <http://support.microsoft.com/kb/323752> (this page was removed shortly after our paper was published in September 2010, but other copies are still available, for example http://www.digitalsupporttech.com/mskb/323/323752_Session_variables_are_lost_if_you_use_FRAMESET_in_Internet_Explorer_6.htm).

86. Nicole Friess, *Add Amazon.com to the List - Class-Action Lawsuit Alleges Data Privacy Violations*, INFORMATION LAW GROUP (Mar. 11, 2011), <http://www.infolawgroup.com/tags/p3p/>.

harm.⁸⁷ The plaintiffs revised their complaint and the judge ruled in June 2012 that the plaintiffs can proceed with one of their claims.⁸⁸

After our paper was published, we noticed that Facebook updated it's P3P compact policy to the invalid policy:

P3P:CP = "Facebook does not have a P3P policy. Learn why here: <http://fb.me/p3p>"

The link in the compact policy goes to a page that explains:⁸⁹

The organization that established P3P, the World Wide Web Consortium, suspended its work on this standard several years ago because most modern web browsers do not fully support P3P. As a result, the P3P standard is now out of date and does not reflect technologies that are currently in use on the web, so most websites currently do not have P3P policies.

It is likely that if Facebook translated their actual privacy policy into a P3P compact policy, its cookies would be blocked by IE when used in a third-party context due to the fact that Facebook does not provide a way to opt-out of tracking and the lack of statement groups in P3P 1.0 compact policies may result in an over statement of their data use. While Facebook is now trying to be more up front about their bogus P3P compact policy through a human-readable statement, P3P compact policies are not intended to be read by humans. The Facebook compact policy is not meaningful to IE, and it circumvents IE's cookie blocking mechanism, on which many consumers rely. A lawsuit has been filed against Facebook in 2012 that alleges that users were harmed by this bogus P3P compact policy.⁹⁰ A lawsuit was also filed against Google, which has a similar P3P compact policy.⁹¹

Amazon took a different approach, and changed its compact policy to a policy that appears to be syntactically valid. However, the corresponding full P3P policy is not valid, and contains the following text:

87. Venkat Balasubramani, *The cookie crumbles for Amazon Privacy Plaintiffs – Del Vecchio v. Amazon*, TECHNOLOGY AND MARKETING LAW BLOG (December 2, 2001), http://blog.ericgoldman.org/archives/2011/12/the_cookie_crum.htm.

88. Wendy Davis, *Judge Rejects Amazon's Bid To Dismiss Privacy Lawsuit*, ONLINE MEDIA DAILY (June 5, 2012), <http://www.mediapost.com/publications/article/176077/judge-rejects-amazons-bid-to-dismiss-privacy-laws.html>.

89. Excerpted from Facebook's Platform for Privacy Preferences (P3P), FACEBOOK, <https://www.facebook.com/help/?topic=p3p> (last visited on Apr. 23, 2012).

90. In Re: Facebook, Inc. Internet Tracking Litigation, No. 12-MD-02314 (N.D. Cal. filed Feb. 8, 2012).

91. Villegas vs. Google Inc., No. 12-CV-00915 (N.D. Cal. filed March 20, 2012).

Because some browsers require a P3P policy, we have created a compact P3P policy that outlines some, but not all, of the details of our privacy practices. The compact policy relates primarily to our use of HTML cookies and personally identifiable information associated with such cookies. However, we have not included a full P3P policy because the binary limitations of the required XML code currently do not fully and adequately express our policies and practices. Instead, we ask that you read our full Privacy Notice at www.amazon.com/privacy.⁹²

While it is not entirely clear what Amazon's concern is, their reference to "binary limitations" suggests that they are uncomfortable selecting from among some of the multiple-choice P3P vocabulary elements. If Amazon's P3P compact policy matches their actual data practices with respect to cookies, then they are no longer circumventing the IE cookie-blocking mechanism. However, they are not fully complying with P3P either.

In February 2012, the discovery that Google was circumventing Safari cookies led to a couple of blog posts on IE cookie circumvention that set off a flurry of media attention.⁹³

The lack of overall P3P compliance demonstrates the ineffectiveness of P3P as a self-regulatory program. After we found that TRUSTe websites with compact policies were just as likely to have errors as the other websites with compact policies that we surveyed, TRUSTe acknowledged that P3P compliance was not part of their routine review process and they did not expect to make it part of their process.⁹⁴ TRUSTe president Fran Maier stated that "P3P irrelevance resulting from barriers to implementation and disregard by consumers encouraged non-compliance."⁹⁵ Indeed, it seems the industry has all but given up on P3P but cannot abandon it completely as long as Microsoft keeps using it as part of their cookie-blocking filter in Internet Explorer.

Our finding that companies are adopting P3P in order to misrepresent their privacy practices to Internet Explorer's cookie

92. Excerpted from AMAZON.COM's full P3P policy, http://www.amazon.com/w3c/p3p_full.xml (last visited Feb. 8, 2011).

93. See Lorrie Cranor, *Internet Explorer privacy protections also being circumvented by Google, Facebook, and many more*, TECHNOLOGY|ACADEMICS|POLICY (February 18, 2012), http://www.techpolicy.com/Cranor_InternetExplorerPrivacyProtectionsBeingCircumvented-by-Google.aspx (February 18, 2012); Dean Hachamovitch, *Google Bypassing User Privacy Settings*, IEBLOG (February 20, 2012), <http://blogs.msdn.com/b/ie/archive/2012/02/20/google-bypassing-user-privacy-settings.aspx>.

94. Comment of TRUSTe president Fran Maier to *TRUSTe's Response to CMU's Findings (Update 2)*, POGOWASRIGHT.ORG (Sept. 14, 2010), <http://www.pogowasright.org/?p=13959&cpage=1#comment-257>.

95. Fran Maier, *More on the Problem with P3P*, TRUSTE BLOG (Sept. 14, 2010), <http://www.truste.com/blog/2010/09/14/more-on-the-problem-with-p3p/>.

blocking feature – one of the most commonly-used tools that consumers have for protecting their online privacy – raises serious concerns. While this would seem to be an area where the US Federal Trade Commission and other regulators within and outside the U.S. could exert their enforcement authority, as of June 2012, no public enforcement actions have been taken based on P3P.

Sarah Spiekermann and I interpret notice and choice as part of the “privacy-by-policy” approach to data protection, which we distinguish from the “privacy-by-architecture” approach. At the extremes, privacy-by-policy relies on trusting companies to say what they do and do what they say (perhaps with the aid of legal enforcement mechanisms), whereas privacy-by-architecture focuses on data minimization so as to physically prevent data misuse. Privacy-by-architecture often involves the use of technology to anonymize data or allow data to be processed locally with minimal or no transfer. In between, hybrid approaches supplement privacy policies with technical enforcement mechanisms.⁹⁶ Since the privacy-by-policy approach does not guarantee data will not be misused, the existence of an effective enforcement mechanism is critical to its success. P3P is an example of a technical mechanism designed to support, but not enforce, the privacy-by-policy approach. As we have seen, since there has been no external enforcement that P3P policies are accurate, P3P has become a useless standard.

VI. OPTING OUT OF ONLINE BEHAVIORAL ADVERTISING

Online behavioral advertising (OBA), defined by the FTC as “the practice of tracking consumers’ activities online to target advertising,”⁹⁷ has been the focus of much of the FTC’s privacy-related attention since 1996. In response to FTC pressure, a group of online behavioral advertising companies launched a self-regulatory organization called the Network Advertising Initiative (NAI) in 1999. The NAI published a set of principles in 2001 and revised them in 2008. As of October 2011, the NAI had 74 member companies listed on its website.⁹⁸ The NAI also runs a consumer opt-out service on its website that offers a central place for consumers “to ‘opt out’ of the behavioral advertising delivered by our member companies”⁹⁹ by setting opt-out cookies in their web browsers.

96. Sarah Spiekermann & Lorrie Faith Cranor, *Engineering Privacy*, 35 IEEE TRANSACTIONS ON SOFTWARE ENGINEERING 67 (2009).

97. *Online Behavioral Advertising: Moving the Discussion Forward to Possible Self-Regulatory Principles*, Federal Trade Commission (Dec. 2007), <http://www.ftc.gov/os/2007/12/P859900stmt.pdf>.

98. The full NAI membership list is available online at NETWORK ADVERTISING INITIATIVE (“NAI”), <http://www.networkadvertising.org/participating/> (last visited Apr. 30, 2012).

99. NAI, *Opt Out of Behavioral Advertising*, <http://www.networkadvertising.org/>

In 2009 the NAI joined with several other industry organizations to form the Digital Advertising Alliance (DAA),¹⁰⁰ which published its own set of principles,¹⁰¹ similar to the NAI principles, in order to demonstrate that the industry could adequately self-regulate. The DAA principles require, among other things, that companies provide a mechanism for opting out of data collection for OBA and that companies provide an “enhanced notice” in the form of a “clear, meaningful, and prominent link” to a disclosure about OBA on every page “where OBA data is collected or used.”¹⁰² The DAA has also introduced a standard “Advertising Option Icon,” shown in Figure 5, to be used next to enhanced notice links. The DAA has engaged the Direct Marketing Association and the Council of Better Business Bureaus to monitor compliance with the principles. In addition, they have approved TRUSTe, DoubleVerify, and Evidon to provide services that assist companies in compliance with the principles.¹⁰³



FIGURE 5: ADVERTISING OPTION ICON SHOWN ABOVE THE TOP RIGHT CORNER OF AN ONLINE AD

When users click on the advertising option icon or accompanying “AdChoices” link they are taken to a page with information about what

managing/opt_out.asp (last visited Apr. 30, 2012).

100. For a list of affiliated organizations see DIGITAL ADVERTISING ALLIANCE (“DAA”), <http://www.aboutads.info/associations> (last visited Apr. 30, 2012).

101. *Self-Regulatory Principles for Online Behavioral Advertising*, DAA (July 2009), <http://www.aboutads.info/resource/download/seven-principles-07-01-09.pdf> (hereinafter “Self-Regulatory Principles”). For more background on the formation of the DAA see Davis & Gilbert LLP, *Newly Formed Digital Advertising Alliance Announces Self-Regulatory Program For Online Behavioral Advertising* (October 2010), http://www.dglaw.com/images_user/newsalerts/AdvMktngPromo_Behavioral-Advertising-Self-Regulatory-program.pdf; see also Press Release, Better Business Bureau, Major Marketing/Media Trade Groups Launch Program to Give Consumers Enhanced Control Over Collection and Use of Web Viewing Data for Online Behavioral Advertising (Oct. 11, 2010), available at <http://www.newyork.bbb.org/article/major-marketing/media-trade-groups-launch-program-to-give-consumers-enhanced-control-over-collection-and-use-of-web-viewing-data-for-online-behavioral-advertising-22618>.

102. Self-Regulatory Principles, *supra* note 101.

103. *Digital Advertising Alliance Begins Enforcing Next Phase of Self-Regulatory Program for Online Behavioral Advertising*, 4 A’s (May 23, 2011), http://www.aaa.org/news/press/Pages/052311_digital_next.aspx.

companies are collecting information in order to provide targeted advertising and an opportunity to opt-out of OBA. The opt-out links may take users to an opt-out page provided by an advertising company or to centralized opt-out pages provided by the DAA, NAI, or one of the approved service providers. Here users can opt-out of targeting from individual companies or from all companies listed on the opt-out page. When a user requests to opt-out, an “opt-out cookie” is set for each company the user opts-out from. This cookie replaces a cookie containing a unique identifier for tracking users. If users delete their cookies, they will likely inadvertently delete the opt-out cookie and nullify the opt-out. Once opted-out, users should receive generic advertisements rather than targeted advertisements from the companies they opted-out of. However, contrary to users’ expectations, they may still be tracked.¹⁰⁴

The Do Not Track header is part of one of the most recent efforts to provide a mechanism that allows users to control OBA. What started out as a simple proposal to add an extra header to web requests that would signal to websites that a user did not want to be tracked led to the creation of a W3C working group¹⁰⁵ that is struggling to reach a consensus on the meaning of tracking (and not tracking) and to standardize Do Not Track. Despite not having a standard, Mozilla implemented Do Not Track in Firefox and Microsoft implemented it in Internet Explorer in 2011. These implementations allow users to turn the Do Not Track feature on and off. When turned on, the header is sent to all websites a user visits. However, few websites currently act on the Do Not Track header.

In February/March and July/August 2011, we reviewed the websites of NAI members and examined ads on the top 100 websites to check for compliance with DAA notice and choice requirements. We also tested the DAA and NAI opt-out websites. While we found that most NAI members were in partial compliance, we still found many instances of non-compliance, generally related to the enhanced notice requirement. We did find an increase in compliance between the spring and summer checks. However, even two years after the DAA published its Self-Regulatory Principles there were still compliance gaps and the compliance monitoring organizations were just starting to address this.¹⁰⁶

104. Pedro G. Leon, Blase Ur, Rebecca Balebako, Lorrie Faith Cranor, Richard Shay, & Yang Wang, *Why Johnny Can't Opt Out: A Usability Evaluation of Tools to Limit Online Behavioral Advertising*, CHI 2012. Also available as CARNEGIE MELLON UNIVERSITY, CYLAB TECHNICAL REPORT, CMU-CYLAB-11-017 (Oct. 31, 2011) http://www.cylab.cmu.edu/files/pdfs/tech_reports/CMUCyLab11017.pdf.

105. W3C TRACKING PROTECTION WORKING GROUP, <http://www.w3.org/2011/tracking-protection/> (last visited Apr. 30, 2012).

106. It wasn't until November 2011 that the BBB announced that they had begun formal

In addition, we found several problems with the opt-out pages themselves, suggesting that there are bugs in the opt-out cookie program that still need to be worked out.¹⁰⁷

In August 2011 we interviewed 48 Internet users from the Pittsburgh area about Internet privacy, OBA, and the tools they could use to control OBA. We found that interview participants had little understanding of OBA and had no familiarity with the Advertising Options icon. When we showed them ads with the icon and tagline, most participants said they would be unlikely to click on the icon. For example, some did not realize the icon was clickable, some thought the icon was intended for advertisers, and some thought clicking on the icon might lead to more advertising. In addition, participants were unfamiliar with most of the online advertising companies and, for the most part, said they would likely make decisions about allowing or blocking tracking based only on their familiarity with each company.¹⁰⁸

We divided participants into nine study conditions and introduced the participants in each condition to a tool designed to limit OBA. Participants in each condition were shown a different tool and asked to install it on a laptop in our lab and configure it according to their personal preferences. We then asked them to perform a series of tasks using the tool they had installed. We tested nine tools across three categories: opt-out tools, browser built-in settings, and blocking tools. The opt-out tools included opt-out websites provided by the DAA and Evidon, as well as a bookmark tool called PrivacyMark that sets opt-out cookies. All of the major web browsers have built-in privacy tools; however, we tested only the privacy tools built into Firefox 5 and Internet Explorer 9, including cookie-blocking and Do Not Track

enforcement of the OBA principles. The BBB identified six companies that had failed to comply with the principles and announced that all six companies had voluntarily taken steps to become compliant. *See Accountability Program Achieves Voluntary Compliance with Online Behavioral Advertising Self-Regulation*, BETTER BUSINESS BUREAU (Nov. 8, 2011) <http://www.bbb.org/us/article/accountability-program-achieves-voluntary-compliance-with-online-behavioral-advertising-self-regulation-30529>.

107. Saranga Komanduri, Richard Shay, Greg Norcie, Blase Ur, & Lorrie Faith Cranor, *AdChoices? Compliance with Online Behavioral Advertising Notice and Choice Requirements*, CARNEGIE MELLON CYLAB (October 2011), http://www.cylab.cmu.edu/research/techreports/2011/tr_cylab11005.html (to be published in I/S in 2012).

108. Blase Ur, Pedro G. Leon, Lorrie Faith Cranor, Richard Shay, & Yang Wang Smart, *Useful, Scary, Creepy: Perceptions of Behavioral Advertising*, *Symposium On Usable Privacy and Security* (July 2012). We also conducted a much larger online survey to examine user reactions to the AdChoices icon and tagline. We found similar results in our online survey: once again, users were unfamiliar with the icon and unlikely to click on it. We compared the AdChoices tagline with several other taglines and found that AdChoices was worse than a number of alternatives. *See* P.G. Leon, J. Cranshaw, L.F. Cranor, J. Graves, M. Hastak, B. Ur. *What Do Online Behavioral Advertising Disclosures Communicate to Users?*, Technical Report CMU-CYLAB-12-008 (April 2, 2012), http://www.cylab.cmu.edu/research/techreports/2012/tr_cylab12008.html

features. The blocking tools prevent tracking by blocking communications to a tracking website. The blocking tools we tested were Ghostery, TACO, Adblock Plus, and Tracking Protection built into IE9. We found significant usability problems with all nine tools. Many of the tools included jargon and were difficult for users to understand and configure. Most users stated that they wanted to block most or all OBA and believed they had configured the tools to do so. However, in many cases they erroneously concluded that the tool they were using was blocking OBA when, in fact, it was not. In addition, users found it challenging to make meaningful decisions about which advertising companies to block because they were unfamiliar with most of these companies.¹⁰⁹

Our study uncovered a variety of usability problems that could likely be addressed if privacy tool creators paid more attention to user interface design and conducted usability evaluations. Many of the problems where users had trouble finding configuration options or had misconceptions about what a tool was doing are likely fixable. However, even if the tools are made more usable, a model in which users need to make decisions about dozens or potentially hundreds of trackers on a case-by-case basis is unlikely to lead to meaningful privacy decisions. On the other hand, an all-or-nothing model in which all tracking is either allowed or blocked is likely to frustrate users when it interferes with desired web functionality. Trackers perform a variety of functions, including functions users may find convenient or desirable. Ultimately, users may want to allow tracking by “good” trackers and disallow tracking by “bad” trackers, but they have little ability to distinguish the two. Indeed, even experts are hard-pressed to determine what individual trackers do and how the data they collect will be used, and there is much debate about what forms of tracking are acceptable or desirable. Tene and Polonetsky suggest that “policy makers and self-regulatory leaders should coalesce around a common approach to the information-for-value business model”¹¹⁰ that tracking supports and determine whether it is beneficial to society. They argue that once a consensus is reached, then default settings could be implemented consistent with this consensus and most users could accept the default settings and not concern themselves with decisions about individual trackers.¹¹¹

The results of our OBA studies suggest that the notice and choice approach has not been all that successful thus far in helping users control OBA. The industry has been slow to implement their own guidelines and

109. Leon et. al., *supra* note 104.

110. Omer Tene & Jules Polonetsky, *To Track or 'Do Not Track': Advancing Transparency and Individual Control in Online Behavioral Advertising* 55 (Aug. 31, 2011), available at SSRN: <http://ssrn.com/abstract=1920505>.

111. *Id.* at 22.

to build the infrastructure necessary to monitor compliance. With the compliance monitoring efforts, approved service providers, and the recently-launched effort in the W3C to develop a Do Not Track standard, there currently seems to be more momentum behind OBA self-regulation than there ever was behind P3P. Even so, we still have not achieved the conditions necessary to provide meaningful control to users or to ensure that this privacy-by-policy approach will be accompanied by powerful enforcement. User empowerment tools suffer from serious usability flaws. Even if those flaws were corrected, it is not clear that the current approach of asking users to distinguish among hundreds of trackers is actually viable. Industry compliance monitoring is backed up by threats to turn over information about compliance gaps to the FTC. However, unless a company misrepresents their privacy practices, it is not clear that the FTC has the authority to enforce compliance. In addition, the FTC is not sufficiently staffed to pursue actions whenever complaints are filed.

VII. CONCLUSIONS

The free market notice-and-choice approach has been the dominant approach to data privacy in the US. The adequacy of this approach has been discussed and revisited repeatedly since at least 1996 when the FTC held its first Internet privacy workshop. In the years that followed we have seen a continuous cycle of new industry initiatives to improve notice and choice mechanisms and empower individuals, followed by a loss of interest in these initiatives when pressure from regulators subsides. Industry organizations, privacy seal programs, privacy enhancing technologies, and even whole privacy-oriented technology companies have come and gone. The Internet is littered with the remains of these past privacy initiatives. Some are absorbed by new initiatives or quietly forgotten. But some, like P3P, survive in a state of limbo where they arguably are doing more harm than good.

We may be nearing the top of the latest wave of privacy interest, with recent reports from the FTC and DoC, and frequent Congressional hearings on privacy. The AdChoices icon and Do Not Track initiatives are the latest self-regulatory approaches. Already the industry has put considerable effort into AdChoices and is developing an infrastructure to monitor and enforce compliance. Despite a lack of agreement on exactly what it means or a formal standard, web browsers are being shipped with Do Not Track features and some websites claim to respect the Do Not Track header. While these are all positive signs, the lack of regulatory enforcement mechanisms puts the success of these initiatives in doubt.

The experience over the past fifteen years demonstrates that privacy user empowerment tools and notice and choice mechanisms are insufficient to protect privacy. However, as many have suggested, these

tools may be quite complementary to privacy regulation.¹¹² Once consumers are confident that their information is protected at least at a baseline level, standardized notice and choice mechanisms have the potential to provide meaningful control over secondary data uses and sharing. But enforcement mechanisms are needed to ensure that users' choices are respected.

Based on my past work and observations I offer the following recommendations and conclusions.

Incentives for adoption and mechanisms for enforcement are essential. We are unlikely to see widespread adoption of a privacy policy standard if we do not address the most significant barrier to adoption: lack of incentives. If a new protocol were built into web browsers, search engines, mobile application platforms, and other tools in a meaningful way such that there was an advantage to adopting the protocol, we would see wider adoption. However, in such a scenario, there would also be significant incentives for companies to game the system and misrepresent their policies, so enforcement would be critical. Incentives could also come in the form of regulations that require adoption or provide a safe harbor to companies that adopt the protocol. Before we go too far down the road of developing new machine-readable privacy notices (whether comprehensive website notices like P3P, icon sets, notices for mobile applications, Do Not Track, or other anything else), it is essential to make sure adequate incentives will be put in place for them to be adopted, and that adequate enforcement mechanisms exist.

Standardization benefits consumers. There seems to be a clear advantage of standardized notices for consumers. Standardized notices facilitate comparisons and allow consumers to become familiar with terminology and where to look to find particular types of information. However, to be effective, standardized notices need to have fairly rigid requirements so that their elements are directly comparable. An earlier attempt at standardized privacy notices, the layered notice developed by The Center for Information Policy Leadership,¹¹³ introduced some good ideas, but allowed so much flexibility that companies ended up using it in fairly inconsistent ways. While there may be a place for companies to customize standardized formats to provide specific details, the overall format needs to be fairly uniform.

Machine-readable policies allow for automation. As online interactions get more complicated, it becomes increasingly difficult for users to understand what parties are involved let alone sort through each

112. COLIN J. BENNETT & CHARLES D. RAAB, *THE GOVERNANCE OF PRIVACY: POLICY INSTRUMENTS IN GLOBAL PERSPECTIVE* 155 (Burlington: Ashgate, 2003); M. Ryan Calo, *Against Notice Skepticism In Privacy (And Elsewhere)*, 87 NOTRE DAME L. REV. 1027 (2012), available at http://www.nd.edu/~ndlrev/archive_public/87ndlr3/Calo.pdf.

113. THE CENTER FOR INFO. POLICY LEADERSHIP, *supra* note 57.

of their privacy policies. Thus machine-readable policies play an important role because they allow web browsers and other tools to consume policy data automatically and take actions on a user's behalf (blocking cookies and other forms of profiling, warning users, etc.). Users would benefit from tools they could set once in their web browser or even for all programs on their computer that would provide high-level privacy controls without requiring users to make decisions on a tracker-by-tracker basis (although users may still see benefit in making exceptions for some trackers). Machine-readable policies also have benefits for business-to-business interactions, because they allow businesses to more easily determine the policies associated with their service providers and advertising agents.

Layers allow for both simple and detailed views. An extremely simple privacy notice, perhaps in the form of an icon, is likely to appeal to most consumers. On the other hand, some consumers and privacy experts will want to see more detailed disclosures, and in some cases detailed disclosures are required for legal purposes. A layered approach to privacy notices would make very simple notices readily available with links to more detailed notices.

Standard policy types could simplify privacy decision-making. One way of distilling complicated privacy policies down to a small number of icons (similar to the Creative Commons approach) is to identify the most important practices that consumers are likely to want to know about and develop a small number of policy templates that incorporate these practices. For example, a type I policy might commit to not collecting sensitive categories of information and not sharing personal data except with a company's agents, while a type II policy might allow collection of sensitive information but still commit to not sharing them, a type III policy might share non-identified information for behavioral advertising, and so on. Companies would choose which policy type to commit to. They could advertise their policy type with an associated standard icon, while also providing a more detailed policy. Users would be able to quickly determine the policy for the companies they interact with. In addition, the establishment of a clear set of policy types would likely encourage companies to improve their privacy practices so that they could associate themselves with a more privacy-friendly policy type.

The P3P vocabulary offers a good starting point for future privacy vocabularies. As I've discussed, the P3P vocabulary has been criticized for its complexity and for its lack of expressiveness. There are clearly some areas that could use some fine-tuning, but after about a decade of use, it seems that overall the P3P vocabulary seems to do a pretty reasonable job. As long as companies are forced to use a fixed vocabulary with multiple-choice fields to express their policies, there are

likely to be complaints that any vocabulary is not expressive enough. I recommend collecting very specific examples of problems companies are having expressing their policies in P3P, and using these to help frame discussions about where changes are needed.

