Do Not Track as a Generative Approach to Web Privacy

Jonathan Mayer[1]

Consider behavioral advertising as a hypothetical negotiation problem.[2] On one side of the table is the average user, who wants to access an advertising-supported service—but only give up some privacy in exchange.[3] On the other side is the average online business, glad to provide a service to the user—if able to display an ad, and preferably an interest-targeted one.[4] In the status quo the user is tracked, and the site delivers an interest-targeted ad: the user gets her least preference, and the site gets its greatest preference.[5] But suppose the site could deliver a privacy-preserving interest-targeted ad. The user would be better off, and the site would be no worse off.[6]

Technologies exist for privacy-preserving interest-targeted advertising—they just haven't been adopted.[7] This paper argues that privacy-friendly advertising and similar gains could be achieved by moving privacy choices to a generative platform, and it shows how Do Not Track will do just that.

---

[1] Ph.D. & J.D. student, Stanford University; Student Fellow, Stanford Center for Internet and Society.

[2] This discussion is greatly simplified for clarity. Some users are accepting of third-party tracking. The hypothetical omits the role of advertising networks, defines the status quo as solely behavioral advertising, and assumes that a site marginally prefers to display a behavioral ad. For an empirical analysis of these issues, see Jonathan Mayer, *Do Not Track Is No Threat to Ad-Supported Businesses*, CENT. FOR INTERNET & SOCIETY (Jan. 20, 2011), http://cyberlaw.stanford.edu/node/6592.

[3] Studies have consistently shown that users overwhelmingly reject third-party web tracking. *See, e.g.*, *E.g.*, Joseph Turow et al., *Americans Reject Tailored Advertising and Three Activities that Enable It* 15 (Sept. 29, 2009), *available at* http://ssrn.com/abstract=1478214; Lymari Morales, *U.S. Internet Users Ready to Limit Online Tracking for Ads*, GALLUP (Dec. 21, 2010), http://www.gallup.com/poll/145337/internet-users-ready-limit-online-tracking-ads.aspx.

[4] *See* Mayer, *supra* note 2.

[5] *See* Julia Angwin, *The Web's New Goldmine: Your Secrets*, WALL ST. J., July 30, 2010.

[6] All else being equal, of course.

[7] *E.g.*, Vincent Toubiana et al., *Adnostic: Privacy Preserving Targeted Advertising*, PROC. 17TH ANN. NETWORK & DISTRIBUTED SYS. SECURITY SYMP. (2010), *available at* http://crypto.stanford.edu/adnostic/adnostic-ndss.pdf; Matthew Fredrikson & Ben Livshits, *RePriv: Re-Envisioning In-Browser Privacy* (Microsoft Research Technical Report MSR-TR-2010-116, 2010), *available at* http://research.microsoft.com/pubs/137038/tr.pdf.

The Platform for Privacy Preferences (P3P)

The notion of a privacy negotiation is nothing new.

The original web suffered from amnesia. Quit your browser and every interactive site was reset. And so, in 1994, a Netscape engineer implemented a fix: the cookie, a remotely accessible data store within the browser.[8]

Just three years later, every major browser supported cookies. Users could save shopping carts; they could store preferences; and they could maintain a login. But users' activities also could be—and and increasingly were—tracked, not only by the sites they visited but also by invisible third parties.

Recognizing the privacy threat, a group of concerned computer scientists began work on the Platform for Privacy Preferences (P3P), a technical mechanism for a privacy negotiation between a user and a website. A user would declare her privacy preferences to her browser, and a site would declare its privacy policy in a computer-interpretable form. Upon visiting a site, the browser would match the user's preferences to the site's policy. If the two aligned, the browser would load the site. If not, the user would have a choice of whether to allow the site anyways or use site-specific, issue-by-issue opt outs.

The protocol specification aimed to be sufficiently fine-grained and flexible to capture the nuance of privacy policies. A site could, for example, indicate it would share a user's ZIP code, pager number, and political affiliations with an advertising network, but keep to itself her age, employer, and health records. Likewise a user could fine-tune privacy preferences, such as allow sites to share purchase history and general interests, but not financial information.

The P3P project intended to release a standard in eighteen months.[9]

---

[8] John Schwartz, *Giving Web a Memory Cost Its Users Privacy*, N.Y. TIMES, Sept. 4, 2001.

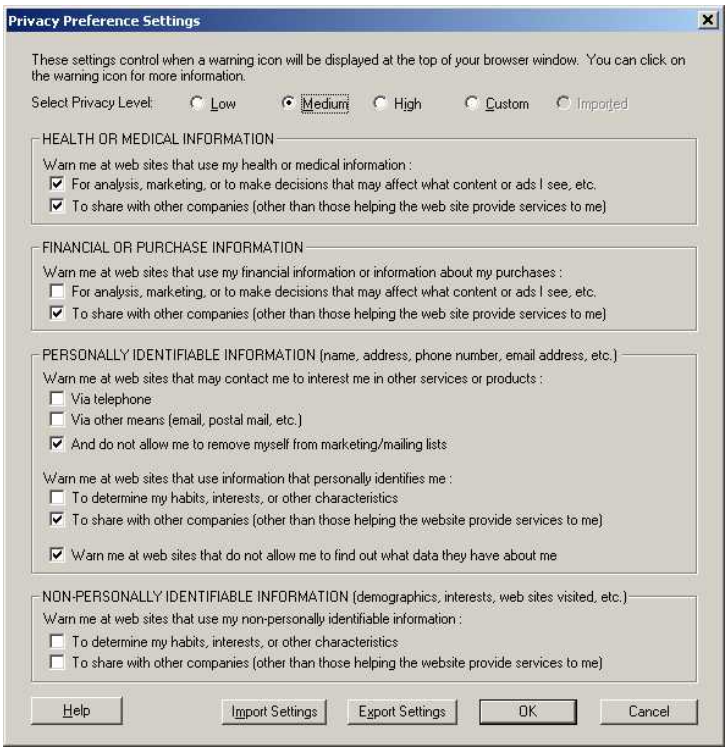

P3P Browser Preferences[10]

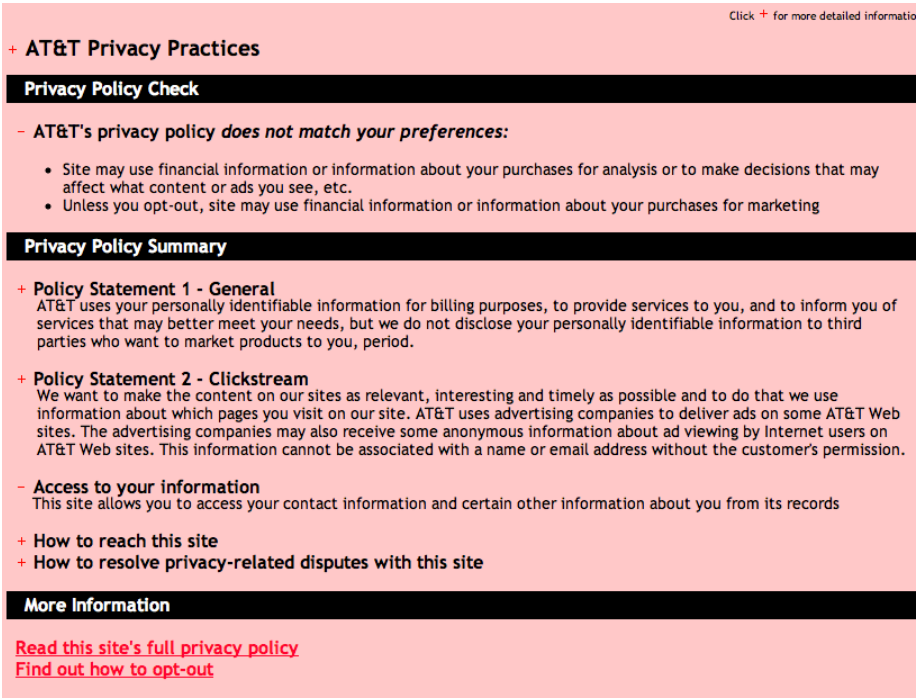

P3P Policy Warning[11]

---

[9] Platform for Privacy Preferences Project, *Project Update* (July 10, 1997), http://www.w3.org/P3P/100797Update.html.

[10] Privacy Bird, *Privacy Bird Tour*, http://www.privacybird.org/tour/1_3_beta/tour.html.

It took five years; P3P was finally standardized in 2002.[12] But few tools existed for creating policies, only a minority of sites adopted P3P, and web browsers implemented only bits and pieces of the standard. After a final effort to reinvigorate the project, in late 2006 the P3P standards group unraveled.[13] Few P3P policies remain, and most do not conform to the standard.[14]

Generativity and Privacy Choice

In the wake of P3P's failure, critics have launched a number of assaults: it presented users with far too many and too complex choices;[15] it was difficult to enforce;[16] and its language was inadequate for capturing the nuance of privacy policies.[17] All fair points. But here's one more, which I view as the most fatal: P3P was not generative.

In *The Future of the Internet—And How to Stop It* Jonathan Zittrain endeavored to identify the properties of technologies that lead to explosive, unguided innovation. He argued for five factors, technologies that[18]

- Make difficult tasks easier;
- Are easily adapted to new purposes;
- Require little to no expertise or training;

---

[11] *Id.*
[12] LORRIE CRANOR ET AL., THE PLATFORM FOR PRIVACY PREFERENCES 1.0 (P3P1.0) SPECIFICATION (Apr. 16, 2002), *available at* http://www.w3.org/TR/P3P/.
[13] LORRIE CRANOR ET AL., THE PLATFORM FOR PRIVACY PREFERENCES 1.1 (P3P1.1) SPECIFICATION (Nov. 13, 2006), *available at* http://www.w3.org/TR/P3P11/.
[14] Pedro Giovanni Leon et al., *Token Attempt: The Misrepresentation of Website Privacy Policies Through the Misuse of P3P Compact Policy Tokens*, PROC. 9TH ANN. ACM WORKSHOP ON PRIVACY IN THE ELECTRONIC SOC'Y (2010).
[15] Ari Schwartz, *Looking Back at P3P: Lessons for the Future* (Nov. 2009), *available at* https://www.cdt.org/files/pdfs/P3P_Retro_Final_0.pdf.
[16] Ruchika Agrawal, *Why is P3P Not a PET?* (2002), http://www.w3.org/2002/p3p-ws/pp/epic.pdf.
[17] Lorrie Faith Cranor, *Incentives for Adoption of Machine-Readable Privacy Notices* (Nov. 5, 2010), http://www.iab.org/about/workshops/privacy/papers/lorrie_cranor.pdf.
[18] JONATHAN ZITTRAIN, THE FUTURE OF THE INTERNET—AND HOW TO STOP IT 71-73 (2008). *See also* James Grimmelmann & Paul Ohm, Book Review, *Dr. Generative or: How I Learned to Stop Worrying and Love the iPhone*, 69 MD. L. REV. 910 (2010).

- Are easy to learn about and acquire; and
- Facilitate transfer of changes.

Zittrain bundled these properties into a solitary adjective: "generative."

For a privacy choice platform to succeed, it must be generative. New websites, web services, web business models, and web technologies are established daily. As a consequence, web privacy considerations are in constant flux. How would an ossified, purpose-built privacy choice mechanism respond to content-sharing sites? Social networking? Social plug-ins such as the Like button? Single sign-on like OpenID? Would web businesses have to retain privacy platform consultants? Would there have to be associations and conferences just for privacy platform experts?

Such would have been P3P's fate, if it had lasted longer. P3P was difficult to implement for a browser or website, narrowly purposed, convoluted, under-documented, and difficult to generalize across sites. It wasn't generative. And so it failed.

Allocative Technologies

Perhaps a generative privacy choice platform could be developed. I have doubts. But here's an alternative approach: Instead of constructing a new generative platform, why not build on an existing one? And, when a problem does not naturally fall to the generative platform, why not use simple mechanisms—for convenience, "allocative technologies"—to relocate the problem there?[19]

[19] This argument suggests a rough technological parallel to Guido Calabresi's "cheapest cost avoider" thesis: allocate a difficult online problem to the most generative system available.

Language signaling is a common allocative technology. Browsers don't include sophisticated translation software. Instead, they signal a user's language preferences, and it's up to foreign sites to develop alternate-language versions using standard web technologies.

Mobile web browsing now relies extensively on allocative technology. Before the iPhone, most mobile device browsers would attempt (unsuccessfully) to adapt websites for easier viewing on a small screen. Recognizing the failure of this approach, Apple launched its mobile browser with an explicit reliance on allocative technology: Apple encouraged websites to build mobile-friendly versions of their sites using standard, generative web technologies. In response to a request from an iPhone, sites were to redirect to their mobile versions. This allocative approach is so successful that every major mobile browser since has adopted it.

Do Not Track as an Allocative Technology for Privacy Choice

Do Not Track is an allocative technology for privacy choice: it relocates the third-party privacy negotiation from the browser, where it has languished since P3P, to the web. In response to a Do Not Track user's request, a web service is free to respond using the standard web technology toolset. It could just deliver its service and an ad without tracking. Or it could ask a user for her interests to deliver a privacy-preserving interest-based ad. Or it could ask for a small payment. It could even refuse to provide service until the user disables Do Not Track.

And there, at last, is the long-sought web privacy negotiation. Do Not Track gives users a veto of the status quo, and allows web services to respond with meaningful privacy choices built on a generative platform.