The Paradox of Privacy Empowerment: The Unintended Consequences of "Do Not Track"

Position paper for W3C Workshop: Do Not Track and Beyond Berkeley, California, November 26-27, 2012

Berin Szoka[1]

The debate over "Do Not Track" offers an excellent microcosm for understanding the larger privacy policy discourse. Arguments for giving users a tool to express their privacy preferences exert enormous rhetorical appeal. Those arguing for versions of DNT that are more restrictive of the collection and use of information about user behavior essentially insist that "We're merely giving users a choice!" Who could possibly be against letting users choose for themselves? Why should anyone else get to choose *for us*—especially companies that seem to be profiting from the ignorance or helplessness of users?

Tools like "Do Not Track" (and "privacy-friendly" interfaces more generally) are usually justified as simply offering users a means of expressing their true preferences. But such choice architectures[2] are anything but neutral: even with the best of intentions and in the name of facilitating user choice, choice architects will produce outcomes that users would not have chosen if they could make fully rational decisions in a frictionless world without transactions costs. This is the essential paradox of user empowerment.

"Privacy advocates" regularly cite opinion polls showing that users demand greater privacy protection—and thus conclude that privacy-friendly choice architectures simply facilitate the true preferences of users. But listening to what consumers *say* they want tells us much less about their preferences than seeing what preferences they *reveal* in the process of making real-world decisions about trade-offs among values. As much as users value privacy, they do not value privacy in isolation or inherently, but relative to other values—including other forms of privacy.

To avoid the paradox of user empowerment to the greatest extent possible, choice architects must understand how their proposed choice architecture will shape real-world outcomes, and the impact that will have on these many competing values. Let us consider the unintended consequences of three contested aspects of DNT:

[1] This position paper draws testimony I gave to the Senate Commerce Committee in June 2012, http://techfreedom.org/node/185

[2] On term "choice architecture" and its inherent non-neutrality, *see generally* Richard H. Thaler University of Chicago, Cass R. Sunstein & John P. Balz, Choice Architecture, April 2010, http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1583509.

1. **Default setting** - How, and by whom, may a browser be set to send DNT:1?
2. **Definition of tracking** - What is it DNT:1 tells servers not to do?
3. **Architecture of negotiation** - How do sites get users who send DNT:1 headers to opt-in to tracking—and to remain opted-in?

Each is a complicated issue. But all three may be understood, to a degree, in terms of the traditional opt-in and opt-out paradigms. DNT:1 is nothing more than a signal sent by the user's browser expressing a preference not to be "tracked," however defined—after which website publishers, advertisers and other data collectors must somehow negotiate with the user to get him or her to "opt back in" (a term actually used in the TPE[3]) to "tracking" (by granting a site or network a "user-granted exception"). If browsers and other user agents may turn on DNT:1 by default, then the adoption rate of DNT will quickly exceed publishers' "maximum acceptable loss threshold." Below that point it makes little practical sense for publishers and advertisers to bother building an architecture of negotiation, because it is more cost-effective to let DNT:1 users free-ride off those allow tracking (either by setting DNT:0 or by not having it set at all).

Put more simply, if browsers are allowed to turn DNT:1 on by default, most users will live in a world where "tracking" is opt-in. This will be a choice made *for*, not *by*, users. But either way, all of the problems of more general "Opt-In Dystopias" described by Nicklas Lundblad and Betsy Masiello would apply once DNT:1 is turned on. They distill their concerns into four categories:

> **Dual cost structure**: Opt-in is necessarily a partially informed decision because users lack experience with the service and value it provides until after opting-in. Potential costs of the opt-in decision loom larger than potential benefits, whereas potential benefits of the opt-out decision loom larger than potential costs.
>
> **Excessive scope**: Under an opt-in regime, the provider has an incentive to exaggerate the scope of what he asks for, while under the opt-out regime the provider has an incentive to allow for feature-by-feature opt-out.
>
> **Desensitisation**: If everyone requires opt-in to use services, users will be desensitised to the choice, resulting in automatic opt-in.
>
> **Balkanisation**: The increase in switching costs presented by opt-in decisions is likely to lead to proliferation of walled gardens.[4]

The problem is that DNT, like any choice architecture, affects not only "demand" (empowering users to choose) but also the "supply" (the choices available to users). The difficulty of obtaining opt-ins (user-granted exceptions) will serve as a barrier to entry, protecting larger, established incumbents against competition from new entrants. This will be true on some level for individual sites: absent dual-cost structure problem, one might think that any site a user

[3] http://www.w3.org/TR/tracking-dnt/#exceptions-principles

[4] N Lundblad and B Masiello, "Opt-in Dystopias", (2010) 7:1 SCRIPTed 155, http://www.law.ed.ac.uk/ahrc/script-ed/vol7-1/lundblad.asp

visits will easily be able to get an opt-in. But obtaining such opt-ins is costly, both for user and for sites, which must implement a mechanism for obtaining user-granted exceptions. Some sites will simply decide not to risk alienating users, and forego potential additional revenue, while other better established sites or sites less subject to competition, will gain a competitive advantage.

But the greater problem lies with web-wide exceptions, opt-ins to data collection by an ad network or other data collector across the web. To be sure, these are essential to making DNT work without breaking business models that depend on third-party ad networks, but they will also necessarily favor certain established players in the data and advertising ecosystem over other, generally smaller players. One might dismiss these competitive effects as the necessary consequence of restructuring an industry that is loathed by many (despite the benefits it confers),[5] but this consolidation would likely be accompanied by a qualitative change in the *kind* of information collected. Once a network obtains a web-wide exception, why *not* collect more data across the web? Why not associate it in a richer profile? As Masiello and Lundblad explain:

> service providers may attempt to maximise data collection in every instance that they are forced to use an opt-in framework; once a user consents to data collection, why not collect as much as possible? And the increased transaction costs associated with opt-in will lead service providers to minimise the number of times they request opt-in consent. In combination these two behaviours are likely to lead to an excessive scope for opt-in agreements. In turn, users will face more complex decisions as they decide whether or not to participate.[6]

Indeed, why not require users to log-in and provide more information about their real identity? Of course, requiring users to go through an account-creation process would likely turn off many users—if only because it took longer than simply clicking on a dialog box that asked about enabling personalized content. But consumers have become quite accustomed to using Single Sign On systems to log into websites with their Facebook, Twitter, Google or Microsoft Live accounts (and so on). It is not difficult to see such networks becoming federated content networks—the new walled gardens so feared by Tim Wu, Jonathan Zittrain and many others. Leaving a website inside one network and going to the other would require granting another web-wide exception to another network. This isn't necessarily bad but if it ultimately means that *more* information is collected about Internet users, DNT will leave many of its advocates sorely disappointed—and it is certainly not a result any user would have chosen.

This perverse potential (but likely) result simply one example of a larger problem: human rationality is bounded; we are simply not capable of weighing the full implications of choices as complicated as those over privacy. This does not mean that user empowerment is not a

[5] See generally, Comments of Berin Szoka, *Privacy Trade-Offs: How Further Regulation Could Diminish Consumer Choice, Raise Prices, Quash Digital Innovation & Curtail Free Speech*, Dec. 7. 2009 http://ftc.gov/os/comments/privacyroundtable/544506-00035.pdf

[6] *Opt-in Dystopias*.

worthy goal; it is (and it is generally preferable to more top-down alternatives such as regulatory prescriptions on the use of data). But it *does* mean we should not pretend that choice architects are not, in fact, making important choices for users in the process of designing choice mechanisms like Do Not Track.

The problems described above will become more acute the more broadly "tracking" is defined, the more users turn on DNT:1, and the more cumbersome negotiation is. Two particular contested issues within the TPWG will significantly aggravate the opt-in dystopias problem:

1. **Default Settings** - Although the TPWG has always rested on the consensus that DNT headers must be set by users not user agents like browsers,[7] Microsoft breached that consensus earlier this year when it announced earlier this year that it would choose *for* users by setting DNT:1 on by default in its new IE10 browser. European regulators have essentially endorsed this position, calling for users to "told about any default setting; and prompted to keep or to change it"—even if that setting is DNT:1, and therefore not compliant with the DNT spec—and insisting that servers must not disregard DNT headers, even when sent by browsers that turn on DNT:1 by default.[8] It remains unclear how this issue will be resolved.
2. **Configuration -** The TPWG co-chairs recently rejected a proposal to clarify that, to "reflect the user's preference," user agents must "require equal effort to configure [DNT]"[9]—prompting the first formal objection filed in the TPWG.[10] Thus, unless this decision is ultimately reversed by the W3C, a user agent need not set DNT:1 by default if doing so proved problematic; it need only design a user interface that will achieve the same result.

Ultimately these concerns are likely to be dismissed by insistence that sites and services will simply negotiate around DNT to reach the same outcome they would have reached anyway. But in the real world (as opposed to a frictionless perfect market), transactions costs often swamp the gains created by transactions such as the negotiation between site and user. The online advertising ecosystem currently works because it generated tiny amounts of value from enormous volumes of transactions. Even the small transactions costs of forcing today's implicit quid pro quo to become explicit could produce dramatically different outcomes. Nor is it clear that negotiation or payments would generate as much revenue as advertising—meaning that rising transactions costs would be borne by publishers, and passed on to users in the form of reduced quality, quantity or innovation, or higher prices (if they can actually charge prices).

[7] "The goal of this protocol is to allow a user to express their personal preference regarding tracking to each server and web application that they communicate with... Key to that notion of expression is that it MUST reflect the user's preference, not the choice of some vendor, institution, or network-imposed mechanism outside the user's control." TPE § 3.

[8] Neelie Kroes, An update on Do Not Track The Centre for European Policy Studies (CEPS)/Brussels, 11 October 2012, http://europa.eu/rapid/press-release_SPEECH-12-716_en.htm

[9] http://lists.w3.org/Archives/Public/public-tracking/2012Sep/0197.html

[10] http://lists.w3.org/Archives/Public/public-tracking/2012Oct/0104.html

Building on Ronald Coase's seminal work on the importance of transactions costs, Harold Demsetz offered the basic insight that continues to guide the law and economics of setting defaults (which economists generally refer to as "property rights"): in a frictionless world, if the initial assignment of rights is inefficient, negotiation will inevitably and costlessly solve the problem; but in the real world, that initial assignment may prove sticky, thus we should not assign rights in ways that are inefficient.[11] Once again, choice mechanisms are not neutral. If, the day before Microsoft announced their decision to set DNT:1 by default, it was true that "majority default DNT is not the world this standard will exist in. DNT is going to be a 10% solution,"[12] and DNT:1 creates the negative unintended consequences described above (among others), why should choice architects not set the initial assignment to the setting that is more likely to be efficient: DNT:1 *off* by default and not privileged when users configure their browser? An argument could be made to the contrary if it could be shown that "tracking" (as defined by the DNT spec) actually lead to real harm, but as yet, no such argument has been substantiated, and the question of harm has repeatedly been sidestepped within the TPWG.

It is understandable, if ironic, that privacy advocates should desire outcomes that could actually reduce privacy and make consumers worse off—because the chain of causation is attenuated and unclear compared to the noble intentions behind restrictive defaults. Nobody wins Nobel Prizes in Economics for explaining things that are completely obvious, and even once they do, it can take decades (or more) for their insights to permeate areas of discourse outside of economics—such as Internet standard-setting.

It is much more understandable what some market players have to gain by joining forces with well-intentioned but short-sighted privacy advocates: competitive advantage. This is simply another example of the well documented alliance of "bootleggers and baptists."[13] Microsoft, in particular, stands to lose little by disrupting the online advertising market, in which it has struggled to compete. It is by no means clear whether a world of high DNT adoption rates would benefit, in relative terms, Microsoft more than Google (or, for that matter, Facebook), but it might well help Microsoft, since it would generally favor large incumbents with directs relationships with users, such as through the browser and OS. And Microsoft would hardly be the first company to wager that it held a losing hand, and that its odds would be better with a fresh deck of cards.

What lies ahead for choice architects "beyond DNT?" The perpetually difficult task of weighing costs and benefits, and attempting to foresee the unpredictable, in shaping users' choices.

[11] Harold Demsetz, *Toward a Theory of Property Rights*, 57:2 Am. Econ. Rev 347 (1967). http://www.econ.ucsb.edu/~tedb/Courses/Ec100C/Readings/Demsetz_Property_Rights.pdf

[12] *See* Lauren Gelman, "Re: tracking-ISSUE-150: DNT conflicts from multiple user agents [Tracking Definitions and Compliance]", public-tracking@w3.org mailing list, May 30, 2012, http://lists.w3.org/Archives/Public/public-tracking/2012May/0341.html.

[13] Bruce Yandle, "Bootleggers and Baptists-The Education of a Regulatory Economist," Regulation 7, no. 3 (1983): 12. http://www.cato.org/pubs/regulation/regv7n3/v7n3-3.pdf