April 29, 2020

Do Not Track, a "handoff"

Nick Doty
UC Berkeley, School of Information

Status of This Document *This is a case for exploration of the Handoffs model for analyzing tech policy and a section of a published dissertation: Enacting Privacy in Internet Standards.*

Do Not Track, a "handoff"

An ad that follows you

Out to lunch with a friend, the conversation drifts to buying holiday presents. You have been struggling recently to come up with an idea for a gift that will surprise your spouse; your friend recommends a particular brand of watch that you haven't heard of before. You pull out your smartphone and type the name into the search box; your friend taps the link to the appropriate online store and shows you a couple of the colors he thinks your spouse might like. Lunch arrives, and you put away your phone and put aside shopping plans for now, there are still a few weeks before the holiday.

That evening, you're sitting on the couch next to your spouse, who mentions a particular news item from the day. Pulling out your laptop, you load an article on the topic and scroll through it; to your shock, you see an ad in the middle of the article for the exact purple watch you were looking at over lunch. Hoping your spouse hasn't seen it, you quickly click "Back" and open another article instead, and see the same ad. "Oh, were you thinking of getting me one of those?" So much for that little surprise.

For days and days afterward, you keep seeing those ads again and again, on your phone, your laptop, the shared tablet that you keep in the kitchen. All the

more frustrating because you've chosen not to get that watch after all, once it wasn't going to be a surprise, but you still see it, in a series of colors, often multiple times in a day. How was it that your phone talked to your laptop, or the watch manufacturer to the different news sites? Who knew you were looking at this particular product and why was that disclosed to your partner? *Who or what was really responsible?*

Could *you* have prevented this scenario? Probably, using existing technology. If you're aware of this problem and thinking ahead of that possible outcome, you might open a "private browsing" tab on your phone before that first search; when you're done looking at different watches, your browsing history is erased (along with associated cookies) and that's probably enough to prevent the "re-targeting" that revealed your shopping plans. Or you could have installed an ad blocker on your web browser at home so that you rarely see ads anyway. Those individual actions may be effective, but is that how we would determine responsibility here? What if the company knew you didn't want to see those ads everywhere and had refrained from showing them? Or could some part of the system have limited the ads so they only popped up on your phone? Could you tell the advertisers not to customize ads in that way or otherwise control what you see?

Handoffs

System overview The Web as a **socio-technical system** is complex in both its makeup and function. Billions of end users use web browsers on personal smartphones, laptops, smart televisions, desktop computers at their local library or Internet cafe. Web sites that those users visit are produced by newspapers, governments, corporations, non-profits, individual hobbyists.

In its simplest conception as a user-operated client requesting a Web site from a single server operated by a host, the parties are clearly separable and easily identified. (See the Web client-server diagram.) But in understanding the typical commercial arrangements used for hosting, caching, analytics, market research and advertising, the picture is more complex. (See the display advertising diagram, for one small portion of that detail.)

This more complicated landscape of interactions can also be made somewhat visible in the system of requests for resources that make up a Web page. What I have found to often be a surprise in presenting the technical architecture of the Web to non-technical audiences, your Web browser typically makes a large number of requests to load all the resources that make a modern, graphically-intensive Web page. That same infrastructure is used for many analytics and advertising-

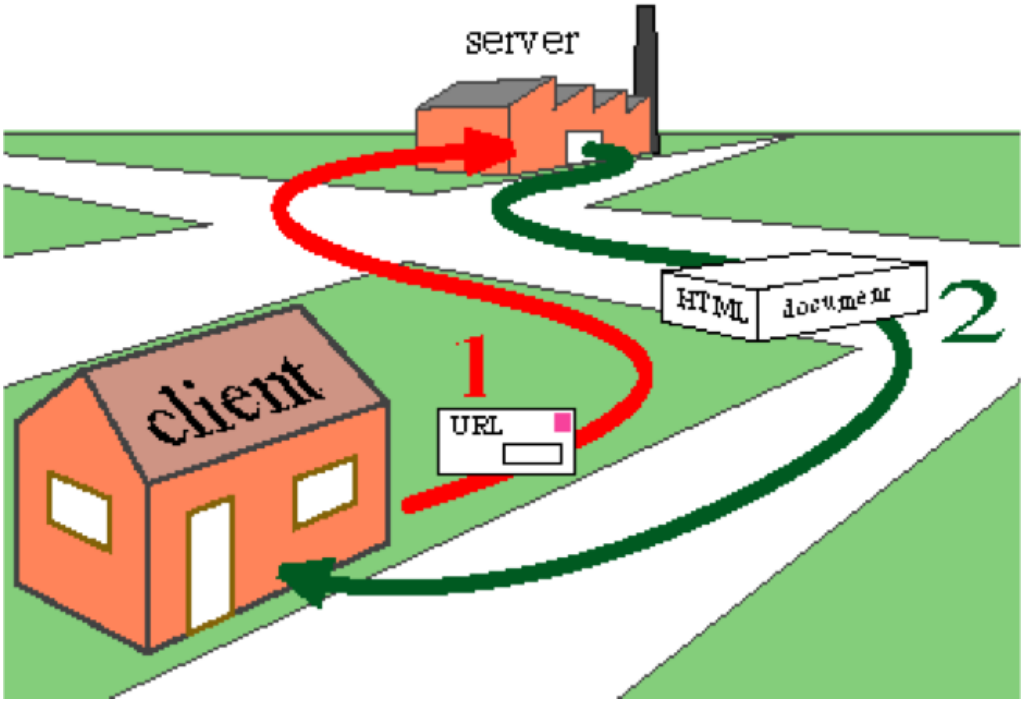

Figure 1: A diagrammatic representation of the Web. Source: CERN.

related purposes; requests are made, behind the scenes, so to speak, to servers operated by analytics and advertising companies, and those communications include information about the user and about the page the user is visiting.

How we define that complicated interconnected socio-technical system and its scope is itself a challenge. Identifying the active stakeholders may be one guide: open multistakeholder processes typically invite participation (or recruit participation) by groups that are likely to be impacted by changes in a particular design. Engagement in political rhetoric or debate also provides an indication of scope. While participants from ISPs were involved in Do Not Track standardization debates, we saw more involvement and focus on the higher layers of the Internet's design; this was a Web topic. Impact on the public, on a larger and less differentiated group of users, of citizens, is harder to gauge this way; nonetheless, consumer advocacy organizations and political figures (including elected officials as well as administrative agency leadership and staff) became deeply involved in Do Not Track efforts.

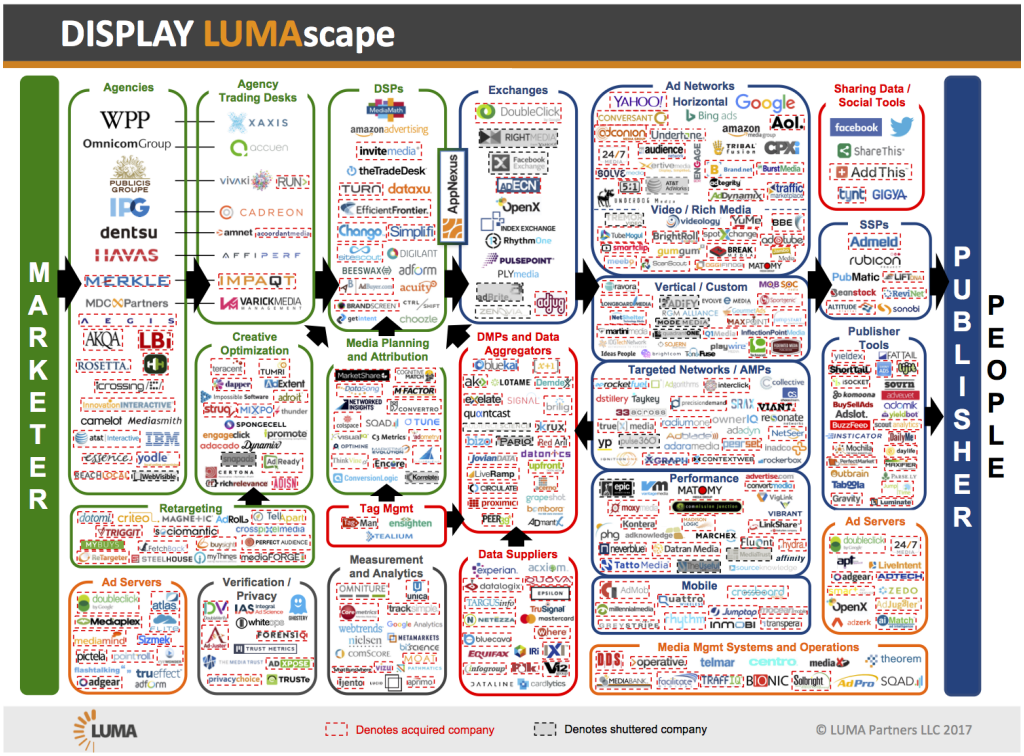

Figure 2: An overview landscape of companies involved in online display advertising; one of a series of popular landscape images from LUMA.

The **actors** that make up our socio-technical system then include both technical pieces (Web browsers, networks, servers), the organizational complexity that arrange those operations (browser developers, advertising networks, analytics vendors), legal and regulatory regimes (the Federal Trade Commission, the EU General Data Protection Regulation), as well as people (users of the Web, individuals who participate in technical standard-setting).

Handoffs between actors Collaborators have defined a **handoff** as the transfer of a function or the responsibility for some value from one actor to another between two different configurations of a system (Mulligan and Nissenbaum 2020). Exploring that shift in responsibility can provide some insight into the political and societal consequences that are too often considered unforeseen or uncontrollable.

Within every configuration of a socio-technical system, there are distributions

of responsibility and functionality – sometimes explicit, but mostly implicit and often misunderstood – among different actors. It can be tempting to think of security in network communications as a value provided purely by technical measures (encryption, say); however, deeper analysis would typically show that security is provided in part by technical measures and in part by legal enforcement, organizational practices, and community norms. In trying to locate responsibility for privacy in our ad re-targeting example, we will come across those rough edges between different actors in the current system, and how the proposed and actual re-configurations of the socio-technical system change how the responsibility for that value is distributed. Understanding why those transfers occur is useful in providing a full explanation of how technological changes affect society.

The history of Do Not Track is so fascinating because we see an attempt to make the distribution of responsibility between technical and legal regulation explicit and because we see an attempt by activists to embody a value in a technical design while explicitly not enforcing that value technologically. These potential handoffs stand in stark contrast to the more unidimensional shifts seen in the high-level trends of automation or privacy-by-design. And seeing this as a handoff better captures the complexity beyond simple comparisons between technical and legal regulation.

Diversity of goals As the functionality of this sociotechnical system depends on the complicated interactions of many different actors, the goals that are implicated for the Web as a sociotechnical system also vary.

Many people use the Web for personal communications: checking their email, posting messages to social networking sites, reading and writing blog posts. Commerce is a common set of uses that is especially relevant to this example: companies provide services for sale; people buy both digital and physical products; online advertising is widespread; media companies provide entertainment services. As we might see more specifically looking at other illustrative examples, there might be very different goals in mind for parties like intelligence agencies or state actors, that may be orthogonal to or in opposition to the goals of many individual users of the system.

We could also identify goals from the stated purposes of designers of the system and its components. The first Web page (Berners-Lee 1992) sets out a succinct and exciting goal for the project:

> The WorldWideWeb (W3) is a wide-area hypermedia information retrieval initiative aiming to give universal access to a large universe

of documents.

Universal access to a large universe of documents gets at the goals of the originators of sharing information, about ongoing scientific projects but also other topics, that can be easily searched and browsed, and implying both retrieval but also easy writing and publication. Berners-Lee even uses the language of the system "aiming to" accomplish that singular goal. However influential that original stated purpose might have been, or might still be among people intimately involved in technical decision-making regarding the Web, it's clear that this system is now complex in a fundamentally different way, that no single person or small group of people has control over the function or the direction. The multistakeholder model of technical standard-setting – through which new functionality for the Web is debated and agreed on – reflects the variety of independent but connected stakeholders that are affected by and jointly implement the Web.[1]

That the socio-technical system does not have a singular, agreed upon goal is useful in understanding the tensions in how to distribute responsibility for a particular function or what values (and what particular interpretation of those values) should be designed for in different configurations.

Paradigmatic changes How can we determine responsibility for providing privacy while browsing the Web, as in our initial motivating example? To illustrate the different distributions of how privacy protection is provided within a system, I describe three different system configurations representing three paradigmatic approaches: first, a cumbersome self-regulatory opt-out regime combined with a set of browser cookie controls; second, a proposed co-operative approach with expressed and respected preferences; and third, an active arms-race of ad and content blocking.

Traditional notice and choice Privacy concerns related to the profiling behind online behavioral advertising have been present as long as that business model has been widespread. In the US, the Federal Trade Commission helped negotiate privacy practices with industry self-regulatory bodies, as part of its initial series of reports and actions on online privacy in the 1990s (Federal Trade Commission 1998; "'Self-Regulation and Privacy Online,' FTC Report to Congress" 1999). The notice and choice model was implemented, in part, through "opt-out cookies" – using the same basic technology (HTTP cookies) typically used for tracking user

[1] For more, see Internet Standard-Setting and Multistakeholder Governance.

activity, an interested user could visit a page in their browser that would set opt-out cookies for each of a potentially large number of online behavioral advertising profilers and that cookie would be sent on subsequent interactions. Promises were made by participating online advertising companies to comply with those self-regulatory codes, including to limit the display of behaviorally-targeted ads. These opt-out cookies have been criticized as cumbersome and ineffective (Dixon 2007; Leon et al. 2012): the process of clearing cookies (which you might do for privacy reasons) would effectively opt the user back into profiling and behavioral advertising; cookies might be set to expire and the participating companies would change over time, so users would need to regularly re-visit and re-install opt-out cookies; and cookies were specific to a single browser, so the same process would need to be applied repeatedly across browsers and across devices; finally, the scope of the privacy choice was unclear or unsatisfying, you might still have your browsing information collected by the same parties using cookies and just not see the targeted advertising until the opt-out cookie expired.

Browsers typically provided a user interface for viewing and clearing cookies, and some experimented with plugins to provide transparency into the different interactions with online services that could track user behavior. But determining which cookies were required for functionality (for account logins and commenting interfaces and shopping carts) and which might be for tracking browsing activity across sites was typically infeasible for the user. User education efforts suggested clearing cookies on some regular basis, but doing so also implied the inconvenience of logging out of sites. Third parties developed browser plugins for blocking trackers, or for blocking the display of advertising altogether. Techniques began to be developed for "re-spawning" cookies; taking advantage of browser bugs, browser plugins or configuration details to maintain identifiers of a user even when cookies were cleared.

In this paradigm, user privacy (at least for the re-targeting example in the anecdote above) is available to the user through cumbersome or uncertain actions on their part, with the legal and normative backing of industry trade associations and a regulatory body, or potentially through technical means, although those means were already being outmaneuvered.

DNT While we might typically identify activists in the area of online privacy as focused on technical solutions, Do Not Track was proposed as a solution that used technology but did not rely on technological enforcement. Rather than continuing an arms race of cookie-management/browser-fingerprinting, an extremely simple

machine-readable signal was to be standardized. Browsers and devices could communicate that signal to other parties online (including both the web sites you visit and the additional parties involved in online advertising and other services), who could comply with the user's expressed preference not to be tracked. Adoption by online parties is voluntary, or at least not enforced by the technical protocols themselves.

In this proposed paradigm, privacy is available as a simple choice to the end user, and that choice is expressed through their browser software and enacted through a similar mix of self-regulatory industry action and the potential for regulatory enforcement. DNT's technical mechanisms are designed specifically to allow for enforcement of a user preference through a combination of consumer regulation, industry self-regulation and software changes. How those choices are enacted, and whether the user understands whether their expressed preference is respected is not technologically enforced, but left up to that combination of private organizational ordering, legal mechanisms and technical designs.

A new arms race Currently, DNT standardization has been completed without widespread adoption by online services and major online advertisers have indicated that they will not modify tracking behavior in response to a user's expressed preference. Industry trade associations and self-regulatory groups have not further developed any alternative browser-based tools. Up to this point, browser vendors have maintained a Do Not Track setting for users, but have also developed more nuanced technical tools for blocking requests or cookies. The use of ad blockers has increased, in add-ons, modes and dedicated browsers. Some publishers rely on vendors to detect ad or tracking blocking and impede or block access to their published content.

While the focus of this analysis has been over distribution of responsibility for the value of privacy, motivated by privacy concerns regarding collection of browser history and disclosure in alternative contexts, this phase of ad-blocking arms race notably involves other values. Ad and tracking blocking software is designed for and advertised as promoting a broader range of values – performance improvements, better security or a less distracting reading experience – in addition to, or instead of, the preservation of privacy.

In this paradigm, competing software design changes – on the client-side and the server-side – impact user privacy, but also security, performance, access to content, and web site business models, with changing implications that are hard for users to measure but may be more visible.

Modes of action In modeling handoffs between configurations, we consider not only the modalities of regulation – markets, law, architecture and norms – used by the various actors within our socio-technical system but also other properties of their actions – whether they are visible or invisible, expressive or coercive – which are described as the **mode** of action.

Of particular relevance here is that we can distinguish between the actions within each of the three paradigms as well as actions used to negotiate or move between those paradigms.

For each paradigm, what are the prominent actors and modes of action and how do they interact?[2]

Modes of action within traditional notice and choice Most prominently featured in the traditional notice and choice paradigm (see Traditional notice and choice, above) are the self-regulatory arrangements: negotiations between the FTC and NAI and certifications and audits of online behavioral advertising organizations. These negotiations are typically private, don't involve direct consumer representation and may be unknown or invisible to the end user.

This opt-out paradigm relies on certain technical arrangements as well. HTTP cookies are re-used for organization-by-organization opt-out communications, and a Web application both explains the opt-out process and allows for setting those opt-out cookies. These are architectural measures that are implemented and controlled by participating online advertising companies, using the existing technology of cookies as it's implemented by Web browsers; the cookies are expressive signals (implementations typically didn't delete other cookies the advertising networks may have set) but the signal is both set and received by the same party. Opt-out cookies are explained and configured through a web page operated by self-regulatory industry groups, rather than a browser setting or control.

The arms race over this tracking activity, especially in leading up to Do Not Track discussions, features different presentations of controls to users by different parties. Browsers provided cookie clearing as a user-initiated method for inhibiting tracking and educational efforts (a kind of norm-setting) suggested clearing cookies as a part of digital hygiene. Optional add-ons for blocking tracking or blocking ads saw some small levels of adoption. Cookie clearing and management sees a technical response in techniques for correlating activity without relying on the persistence of HTTP cookies, including browser fingerprinting and cookie

[2] We could also organize these by the modality of regulation – markets, law, architecture and norms – as I've done in the Encrypting the Web handoff discussion.

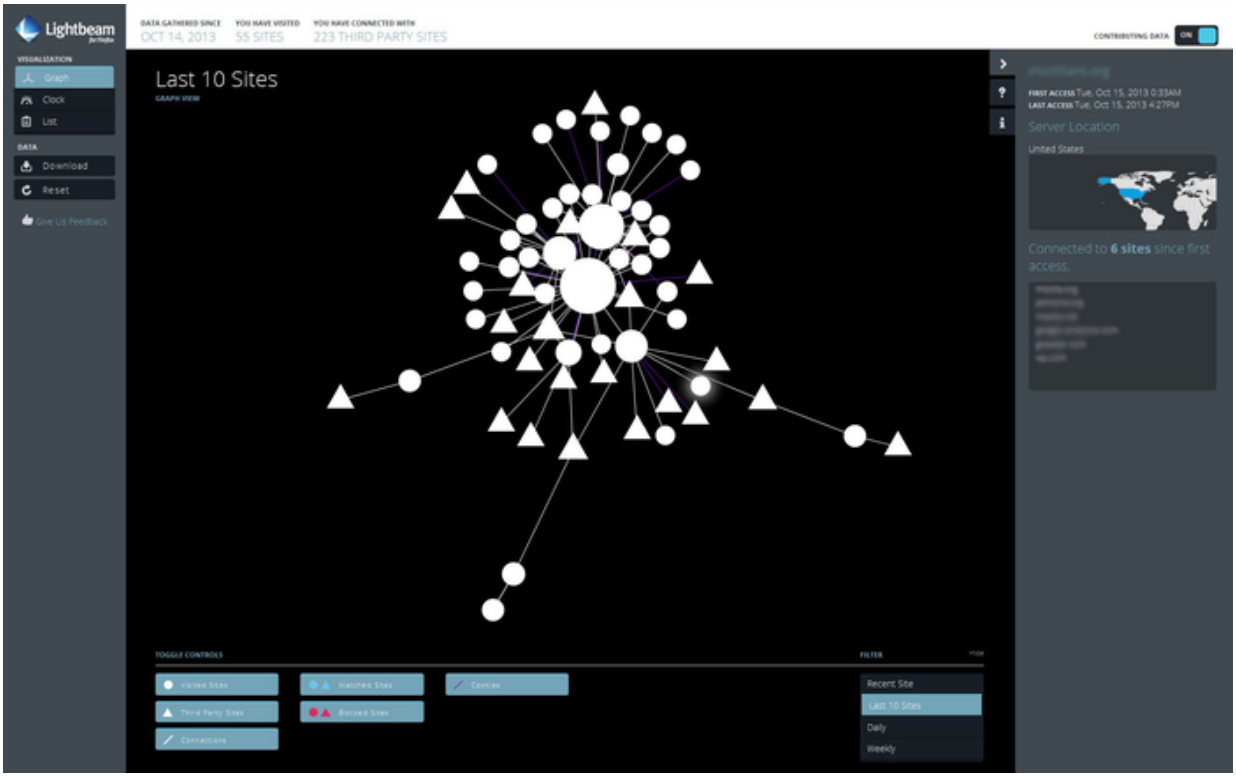

Figure 3: The Lightbeam (previously "Collusion") plugin visualizes common third-party connections from visiting multiple sites.

respawning. While user controls have a direct effect (deleting records stored on their local devices), the arms race makes the effects increasingly obscure and uncertain.

Many technical measures are not self-enforcing mechanisms. Some tools provide increased transparency (including the Lightbeam plugin, pictured) about tracking connections between sites, or the numbers of trackers present. That's an architectural modality of regulation, but it works primarily to persuade or influence other actors, whether it's end users, businesses or regulators.

Modes of action for DNT proposals Do Not Track combines some of the properties of opt-out cookies and direct blocking tools. A DNT header is expressive rather than coercive or self-enforcing: it merely communicates to some other party that a user prefers not to be tracked. But it's also a communication mediated in a different

way than a trade-association-managed opt-out cookie: users have the option to select DNT in their choice of browser software.

DNT as proposed relied on negotiations, if not formal agreements. The standardization process attempts to find consensus among the different parties that might use the DNT header about its meaning and how to comply with it. The W3C standard-setting process is open to a larger and wider variety of stakeholders and its discussions are publicly archived, but this is still largely invisible to the end user.

Enforcement of DNT could happen through distinct means: legal requirements may require or incentivize complying with user preferences in some jurisdictions; statements of compliance may be enforced through trade regulations (for example, FTC enforcement); self-regulatory groups could provide industry agreements and trade associations or other groups could provide external audits of those commitments. Some proposed tying blocking measures to assertions of DNT compliance: tools that block cookies or other tracking mechanisms could refrain from those blocking measures for parties that respond to an expressed preference. That may be a real-time negotiation on behalf of the user ("I'll let you collect some data, so long as you promise to respect my preference not to combine data about me on different sites"), but mediated through expressive signals sent by an online service and client-side measures to block cookies.

Modes of action in blocking and counter-blocking As an implicit or explicit response to the delays in standardization or the lack of server-side adoption of Do Not Track, browser developers have integrated more sophisticated technical responses to tracking. Attempts have been made to systematically block or limit storage while minimizing breakage of popular embedded functionality.[3] Machine learning and other heuristics are increasingly used, beyond the simpler and more static allow and deny lists that were previously proposed. Heuristic, learned and list-based approaches are less direct in the sense that a user-facing control has more complex implications, but the semantic description and the likely implications may at the same time be more comprehensible. "Block tracking scripts" both implies something more complex but also more accessible than deleting a cookie from a particular origin.

As publishers (especially news organizations) increasingly employ paywalls – limits to the number or selection of articles that are available before a user is

[3] For example: Firefox's Tracking Protection and Safari's Intelligent Tracking Prevention and Storage Access API.

Figure 4: A billboard in New York City advertises the Firefox browser based on outwitting online tracking, November 2016.

prompted or required to subscribe – there has also been an increase in blocking access to content for users who are detected as blocking online advertising or tracking. This uses both technical and market measures: the blocking can be accomplished technically, but using pay subscription as an alternative provides a financial incentive to allow advertising and tracking. Market incentives also apply to the browser vendors: performance and privacy protection can be selling points in the competition for users, while the possibility of sites blocking access with a particular browser could cause users to switch.

While the technical means in the DNT paradigm are expressive, blocking and counter-blocking attempt primarily to be self-enforcing or directly effective. The visibility and transparency of these actions is also different: blocking technology can be obscure or opaque (in much the way that tracking technology long has been); paywall prompts are a more explicit, expressive message, and issued from

the publisher to the user rather than from the third-party ad networks. The effects of this shift in responsibility are discussed further below.

Actions that influence the movement between paradigms The previous sections identify the actors and properties of their actions within each of three possible paradigms of our socio-technical system. But those configurations don't exist in parallel or come into being deterministically. We can also observe the actions taken to influence the handoff between different configurations of a system, involving many of the same actors and a diversity of modalities of regulation and modes of action.

One prominent potential starting point in the timeline for Do Not Track is a report from the Federal Trade Commission staff recommending development of a standardized Do Not Track mechanism. This is a notable instance of a government agency actor not using law or rules as its modality of regulation, but rather using communication as a form of norm-setting. Throughout the DNT process, FTC has used diplomacy and encouragement of stakeholder participation, rather than rule-making or bringing enforcement actions.[4]

Participants describe Do Not Track debates as an especially political process, both inside and outside "the room." Lobbying and other kinds of influencing are about setting or changing norms through direct or directed communications. That can involve closed-door lobbying of government officials, certainly, but also public messaging, aimed at users, at companies in the industries involved, or at administrative or legislative representatives. Participants cite references to emails/videos regarding interpretations of a chair's comment at a particular TPWG meeting and a campaign to tie targeted advertising to saving kidnapped children.

Technical and architectural measures are used as means of influencing discussions. Consider two software *patches*[5] introduced during particular moments in DNT standardization debates: a proposed change to Firefox's cookie-setting policy to accept cookies only from visited sites; and a proposed change to Apache's default configuration file to ignore DNT headers sent by Microsoft Internet Explorer. Ultimately, neither of these patches was accepted by the corresponding

[4] The FTC's choice of regulatory actions depends in part on statutory restrictions, historical limitations of administrative rule-making and an approach of engagement, topics covered in great detail by other scholars (Hoofnagle 2016; Bamberger and Mulligan 2015).

[5] A patch is a self-contained proposed change to a piece of software code and is the typical method for introducing, discussing and adopting new changes to collaboratively developed software. The name comes from the older practice of patching over punch cards or paper tape to change a piece of software that was already distributed.

open source software project, but the demonstration of the technical approach was an attempt to influence market forces. While these may not be unique in the history of software development, persuasive software patches are certainly idiosyncratic.[6] This form of communication is also limited in its accessibility: it requires programming expertise, technical reputation or both to contribute these changes, and indeed it takes some technical expertise and understanding of open source software development methods to understand (or translate) the implications of such changes.

Using handoffs

What do we gain from the handoffs model of analysis for the different Do Not Track configurations? In identifying the complex set of actors at different scales; their choice and the mode of their actions; and, the variety of shifts in responsibility that are considered, we can see what is distinctive about Do Not Track and the debate over user privacy of Web browsing activity.

A network of actors and actions Analyzing the socio-technical system as a network of actors and their use of different modalities of regulation can uncover the potentially complicated tensions between various forces at play. This kind of analysis is more familiar in tech policy and science and technology studies as in Actor-Network Theory (Latour 2007) and code-is-law (Lessig 1999). This is just a first step in describing handoffs, but identifying the actors and modal properties of their actions – hard or soft, expressive or self-enforcing, transparent or opaque – can make the implications more explicit for analysis.

An in-depth understanding of Web architecture shows not just the endpoints (the abstract client and server) but also parties that are, abstractly, in the middle, or lower-down: the Internet Service Provider used for connectivity by both the user and the online service; middlebox vendors providing services within enterprises or on in-home networks; the different companies involved in developing and maintaining the user's device, operating system, Web browser and DNS resolution; the parties involved in delivering the diversity of Web pages and their embedded services, analytics, advertising, behavioral tracking and content delivery. Companies are not easily separable into those categories, most notably because many

[6] Another example might be the development of plans for 3-d printed firearms: while some might try to develop and use such weapons, it's commonly accepted that their promotion is an attempt to discourage gun control regulation (Manjoo 2013).

large technology companies compete in multiple areas: Apple sells hardware as well as developing operating systems and a Web browser; Google has the most popular Web browser but the vast majority of its revenue comes from online advertising. Even within the category of online advertising there is diversity of positions: there are different sizes of online advertising networks and different services that different companies provide, and those ad networks and ad technology vendors are distinct from the advertiser itself, that is, a company that has paid in order to show a text or graphical ad for their product or service.[7] That complex network of organizations makes it harder to identify the "sides" in a debate – or even who can speak for or adequately represent what group – or create a simple mapping of who wants what or where a compromise might be. Browser vendors and online publishers might seem like natural mediating parties: browsers might have a closer connection to users and publishers typically have legal agreements and technical measures in place with embedded third parties providing advertising, analytics and other services, but the level of visibility and control that each has is unclear, and our paradigms haven't previously put responsibility on those companies.

It can be tempting to identify categories of technology with the large companies that sell or operate those systems, but in fact there are individual humans who develop software while employed by Google and individual humans who attend meetings with the FTC or visit congressional offices. There may be studies where identifying the individual backgrounds and experiences does not add significantly to an economic analysis of the market positions of the employing organizations, but this is not such an area. Particularly in the Internet field, individuals move between companies and take their experiences and positions with them. Individuals also have multiple roles beyond just their primary professional employment, including their roles in open source software projects and in technical standard-setting bodies. In DNT discussions, roles within companies (engineering vs. sales or product, say) mark a distinct grouping separate from and sometimes orthogonal to employing organization or industry.

This example demonstrates not just a diversity of actors, but the somewhat unusual actions (which vary in their modality of regulation and other modal differences) from our cast of players. In our timeline, the Federal Trade Commission is

[7] It's interesting in this DNT and online privacy context that people who refer to "advertisers" often mean those who sell advertising, like Google and its AdSense network, and not organizations that buy ads, like Coca Cola or car companies, say. Consider the difference between Clear Channel, which might own the large billboard down the street, and Nike, whose ad featuring Colin Kaepernick you might have seen on that billboard. Increasingly, tech companies like Apple and Netflix, are also prominent buyers of outdoor advertising like those billboards.

prominently cited, but not for taking an enforcement action or proposing rules, but recommending a technical mechanism and encouraging standards development. Consumer advocates engage not so much in political lobbying, but join in the technical standard-setting process and provide technical expertise and proposals. Members of Congress send a letter of comments to the World Wide Web Consortium on a public mailing list. Microsoft, a developer of operating systems, a popular Web browser and engaged in online advertising and online publications, makes a prominent default setting proposal. Advertising trade associations are perhaps more conventional in engaging in political lobbying, but perhaps novel terrain in public relations criticisms of non-profits or Web browser businesses.

Shifts in responsibility Specific to handoffs, describing the movement and distribution of responsibility can better explain the impact of decisions and changes that might otherwise be seen as value-free. In this case, we are considering how responsibility for privacy over how data about a user's browsing is collected, shared and disclosed and how that responsibility might be redistributed. The movement between the different paradigms might not be confused for value-free, given the controversy or impact of the different configurations. But the shifts of capability and responsibility are significant and perhaps distinctive in the arena of tech policy. The traditional notice and choice paradigm leaves responsibility unallocated: neither technical guarantees nor regulated arrangements provide a particular sense of confidence about a value like privacy. Instead, as noted in the opening vignette, the end user could execute control[8] if they implemented a set of uncommon technical changes or abstained from using the Web altogether. One response to such a situation of identified inadequate privacy or security protection is to move the discretionary capability away from (or take the burden off of) the end user and instead to provide a technical assurance: for example, a technical system that blocked all data collection that could be used for profiling and behaviorally-targeted advertising. Another response is to set a norm (perhaps bolstered by law, rules or self-regulatory arrangements) for some backstage actors to provide enough of an assurance to the user that they don't need to be concerned with a technical

[8] This example does not speak to the "notice" part of "notice and choice." I don't know that any user has any such capability to understand the technical means behind how ads are tracked and displayed; I've never seen a user successfully use self-regulatory notice icons for that purpose, for example; meanwhile, rumors about how behavioral tracking works and are basically incontrovertible, as anyone knows who has tried to explain to their friends that smartphone microphones aren't constantly listening to their in-person conversations in order to later target an ad for display on Instagram.

arrangement that they don't understand or can't control: for example, laws, rules and self-regulation could prohibit retention of user browsing data or its use for targeting advertising.

Our story here differs from these typical paths. Advocacy and regulatory actors called for a technical mechanism, but not for technical mechanisms that provide guarantees, automatic enforcement and a human-free assurance. Instead, DNT is a technical mechanism for communication of user preferences, rather than traditional notice about business activities, between the user and a subset of other parties. This maintains the opt-out metaphor preferred by businesses and some US policymakers, but with some fundamental differences. Browsers present the choice and information about it to the end user, and can do so in a variety of ways, and users have a new method for communicating with those embedded and often invisible third parties. This is a handoff, but not one that simply removes both capability and assumed responsibility from the end user: instead, it increases communication and makes a kind of shared sense of responsibility between users, browsers (also known as user agents) and the plethora of analytics, advertising and tracking partners.

The new blocking arms race is perhaps more analogous to the security/encryption case. There is still a new handoff, a shift in responsibility: browsers are taking a more direct role in blocking trackers, ads or other resources. These new approaches are less mechanical and less user-directed than the less-widespread alternatives discussed for previous paradigms: there's mostly not a direct list, or a choice of parties to block or unblock, and settings are more likely to be automatic or tied to some other mode rather than user-initiated. Instead, browser developers provide data and algorithms for ongoing identification of tracking and blocking in ways that aren't anticipated to interfere with user-desired functionality. The resulting arms race situation does increase the visibility of the situation for the user, in the case of paywall notices described within the main content of a Web page, and requests for users to provide data explicitly, or become paying subscribers, or to change their browser mode or preferences. Whether and how this situation benefits or diminishes privacy depends on how we conceive of that value. Users of these blocking tools might have less data collected about them but there's little predictability about what tracking is happening when as the different parties try to work around each other's tools. Explicit negotiation with sites over privacy and payment was one of the intended outcomes of Do Not Track as an opt-out mechanism: it makes those tradeoffs more apparent to the user, but might also contribute to different parties collecting different user data (like billing details).

Distinctiveness in handoffs

In our initial example, we saw responsibility for privacy as amorphous and uncertainly placed: who's responsible for this ad that follows you and what can be done about it? By considering different paradigms and the diverse, distinctive actions within and between them, we can evaluate different handoffs of that responsibility between a complex network of actors. Each paradigm – notice and choice, Do Not Track, blocking and counter-blocking – has distinct implications for the value of privacy: whether users have control or rely on others and whether those controls are accessible, effective and enforced technically or through some combination of policies.

The handoff model also helps us analyze the particular properties of the actors and actions within and between those configurations. Debates over DNT included software patches that were effectively persuasive rather than architectural. And Do Not Track is distinctive in being a proposal for a technical mechanism to support user privacy that is expressive rather than self-enforcing and a system that relies on broad multi-party cooperation.

References

Bamberger, Kenneth A., and Deirdre K. Mulligan. 2015. *Privacy on the Ground: Driving Corporate Behavior in the United States and Europe*. MIT Press.

Berners-Lee, Tim. 1992. "The World Wide Web Project." November 3, 1992. http://info.cern.ch/hypertext/WWW/TheProject.html.

Dixon, Pam. 2007. "The Network Advertising Initiative: Failing at Consumer Protection and at Self-Regulation." World Privacy Forum. http://www.worldprivacyforum.org/wp-content/uploads/2007/11/WPF_NAI_report_Nov2_2007fs.pdf.

Federal Trade Commission. 1998. "Privacy Online: A Report to Congress." https://www.ftc.gov/reports/privacy-online-report-congress.

Hoofnagle, Chris Jay. 2016. *Federal Trade Commission Privacy Law and Policy*. Cambridge University Press.

Latour, Bruno. 2007. *Reassembling the Social: An Introduction to Actor-Network-Theory*. Oxford University Press, USA.

Leon, Pedro, Blase Ur, Richard Shay, Yang Wang, Rebecca Balebako, and Lorrie Cranor. 2012. "Why Johnny Can'T Opt Out: A Usability Evaluation of Tools to Limit Online Behavioral Advertising." In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 589–98. CHI '12. New York, NY, USA: ACM. https://doi.org/10.1145/2207676.2207759.

Lessig, Lawrence. 1999. *Code and Other Laws of Cyberspace*. Basic Books. http://books.google.com/books?id=0l1qLyT88XEC.

Manjoo, Farhad. 2013. "The Dumb Argument That 3-D Printers Will Make Gun Control Futile." Slate Magazine. May 8, 2013. https://slate.com/technology/2013/05/3-d-printed-gun-yes-it-will-be-possible-to-make-weapons-with-3-d-printers-no-that-doesnt-make-gun-control-futile.html.

Mulligan, Deirdre K, and Helen Nissenbaum. 2020. "Handoffs." *In Progress*.

"'Self-Regulation and Privacy Online,' FTC Report to Congress." 1999. Federal Trade Commission. July 13, 1999. https://www.ftc.gov/news-events/press-releases/1999/07/self-regulation-and-privacy-online-ftc-report-congress.