TECH FREEDOM

Comments of
Berin Szoka, President
TechFreedom

on

"Protecting Consumer Privacy in an Era of Rapid Change: A Proposed Framework for Businesses and Policymakers"

A Preliminary FTC Staff Report of the
Bureau of Consumer Protection,
Federal Trade Commission

February 18, 2011

Federal Trade Commission Chairman Jon Leibowitz has made privacy the signature issue of his Chairmanship. With his seven-year term on the Commission ending this September,[1] it is understandable that he should feel a sense of urgency to establish a clear legacy in this area by publishing a final version of this preliminary Staff Report[2] before he leaves office—or to help bolster his case that President Obama should re-nominate him, and the Senate should re-confirm him, for a second term on the FTC, so he can stay on as Chairman. Some might blush to speak of such things in agency filings, but there is no shame in acknowledging this reality, and doing so need not impugn the motives of the Chairman or the many dedicated FTC staffers who have worked so hard for so long on this Report.

Indeed, there is much to praise in the FTC Staff Report: Sections II-IV provide an invaluable survey of the history of privacy regulation in the U.S. and the state of the recent debate over privacy in the non-governmental sector, while the "Proposed Framework" in Section V does a commendable job of outlining, as Commissioner Rosch puts it in his separate statement, "a number of 'best practices' that private firms should adopt from the get-go in order to protect privacy."[3] This report has great value in outlining how "Privacy by Design" can actually be implemented by companies to improve privacy practices, both independently and in conjunction with broader self-regulatory efforts.

---

1 *See* Federal Trade Commission, Jon Leibowitz, Chairman, http://www.ftc.gov/commissioners/leibowitz/index.shtml (last modified Feb. 17, 2011).

2 Federal Trade Commission, *Protecting Consumer Privacy in an Era of Rapid Change: A Proposed Framework for Businesses and Policymakers*, Preliminary FTC Staff Report, Dec. 2010, http://www.ftc.gov/os/2010/12/101201privacyreport.pdf (hereinafter "Staff Report").

3 Staff Report, *supra* note 2, at E-2-3, n. 4 (citing report at v, 39, 40-41, 43-52).

I am particularly pleased that FTC staff has not, as many regulatory advocates proposed at FTC Privacy Roundtables and in written comments, abandoned the concept of opt-out in favor of highly restrictive opt-ins for the collection and use of data about consumers. Indeed, whatever else may be said about the "Do Not Track" mechanism endorsed in the Staff Report, it is, in principle, an affirmation of the argument made by defenders of opt-out that enhancing user choice through technological innovation is superior to imposing restrictive defaults.

For these reasons, the Staff Report could make a fine legacy for any FTC Chairman—one that could earn him plaudits from many corners. But in other respects, the Report raises cause for concern. These comments elaborate on the following concerns:

1. **Regulation v. Best Practices**. As Commissioner Rosch notes, however desirable the best privacy practices outlined by the Report might be, "that does not mean that firms should be mandated *de jure* (*i.e.*, by legislation) to adopt them or that firms should be required to do so *de facto* (*i.e.*, that large, well-entrenched firms engaging in "self-regulation" should dictate what the privacy practices of their competitors should be)."[4]
2. **FIPPS**. In particular, the Fair Information Practice Principles (FIPPS) may offer a fine conceptual framework by which businesses can protect the privacy of their users, but they were developed to limit government access to particularly sensitive data (about health) and are therefore not an appropriate framework for dealing with the trade-offs inherent in regulating privacy in general.
3. **The FTC's Authority**. The FTC has not made a clear case that its existing statutory authority to punish unfair and deceptive trade practices is inadequate to protect consumers. The FTC should, as Commissioner Rosch urges, make fuller use of its existing authority. If the agency requires more resources to use that authority effectively, it should request additional appropriations from Congress before seeking more additional powers.
4. **Regulatory Capture**. The FTC must recognize that its interventions in the market, however well intentioned, will always be subject to "capture" by incumbents as weapons against their competitors.
5. **"Do Not Track**." A "Do Not Track" mechanism could, in principle, be an excellent example of how better user empowerment tools can enhance consumer sovereignty and thus decrease the need for paternalistic interventions. Yet once again, it does not automatically follow that government should mandate the use or design of such a mechanism. Technical mandates for "Do Not Track" would, especially at this early stage, amount to having government design the "market for privacy." It would be better for policymakers to let this tool continue to evolve—and let a marketplace between privacy-sensitive users and publishers emerge. The FTC should, however, use its existing authority to hold companies to their promises to respect "Do Not Track."
6. **Costs & Trade-Offs**. Before issuing a final report, the FTC needs much better data about the economic consequences of its proposals in terms of revenue for ad-supported

[4] *Id*.

media and how that revenue is distributed, potential costs to innovation, and the broader competitive landscape of the Internet ecosystem.

7. **Free Speech**. The consequences of regulating the data-based Internet ecosystem are measured not merely in dollars and cents, or in lost innovation, but in terms of expression, speech, media and journalism. Yet the First Amendment has been almost entirely absent from these discussions.
8. **The Rush**. Most of all, I worry that these and other concerns raised in this proceeding, as well as in the comments on the Commerce Department's Privacy Green Paper,[5] cannot be given the attention they deserve between now and September. The FTC should, in general, refrain from calling for increased regulation or new grants of statutory authority in the Final Report. Future arguments for new powers should be made only after addressing the concerns expressed above.

Some will, no doubt, dismiss these concerns as stalling tactics. Yet this would be as unfair as it would be for those concerned about the implications of regulation to dismiss the desire for enhanced consumer sovereignty by refusing to engage in a serious conversation about enhanced choice mechanisms like "Do Not Track."

Instead, my concerns are grounded in a firm belief that sound policymaking can be reduced to a single question: *"And then what?"* What do we imagine will the first order consequences of the various changes the FTC is proposing companies make—or perhaps be required by law to make—be to the Internet ecosystem? If the purpose of a "Do Not Track" mechanism is to create a market for privacy users to essentially, but simply and seamlessly, negotiate with websites over how to fund content, how do we imagine that marketplace will work? Indeed, how would that marketplace evolve under the more elaborate user choice mechanisms recently released by Microsoft or called for by the FTC?

These three, deceptively simple words—"And then what?"—make much the same point the Nobel Prize-winning economist F.A. Hayek made when he remarked in *The Fatal Conceit*, his damning treatment of top-down government planning, that "[t]he curious task of economics is to demonstrate to men how little they really know about what they imagine they can design."[6]

So, how much do we really know about the framework for governing data use the FTC has outlined in its Staff Report? What will be its costs, its effects on competition, its various other unintended consequences? I detail some of these specific concerns below, but readers will find many other concerns more ably expressed in comments filed in this proceeding by those with what Hayek would have called the best "local knowledge"—the technical experts (generally at companies) who are closest to these details.

---

[5] Department of Commerce Internet Policy Task Force, Commercial Data Privacy and Innovation in the Internet Economy: A Dynamic Policy Framework, Dec. 16, 2010, http://www.ntia.doc.gov/reports/2010/IPTF_Privacy_GreenPaper_12162010.pdf (hereinafter "Green Paper").

[6] F.A. HAYEK, THE FATAL CONCEIT: THE ERRORS OF SOCIALISM (W. W. Bartley III, ed. 1988)

I. The Harm Standard & the FTC's Framework

"Fools rush in where Angels fear to tread."[7] But Commissioner Rosch isn't one to rush: He has been a reliable voice of caution on the Commission, who is willing to embrace change (for example, the "Do Not Track" mechanism)—but not recklessly. His separate statement hits the nail on the head: Most of what the Staff Report recommends as "best practices" are, indeed, desirable—"But that does not mean that firms should be mandated *de jure* (*i.e.*, by legislation) to adopt them or that firms should be required to do so *de facto* (*i.e.*, that large, well-entrenched firms engaging in 'self-regulation' should dictate what the privacy practices of their competitors should be)."[8] His explanation of the adequacy and flexibility of the FTC's existing framework bears repeating here:

> As a guide to Congress about what privacy protection law should look like, the Report is flawed. First, insofar as the Report suggests that a new framework for consumer privacy should replace "notice" (or "harm") as the basis for Commission challenges relating to consumer privacy protection, that is unnecessary. ***A privacy notice that is opaque or fails to disclose material facts (such as the fact that consumer information may be shared with third parties) is deceptive under Section 5*** [of the Federal Trade Commission Act]. That is particularly true if the sharing of the information may cause tangible harm. Moreover, Section 5 liability could not be avoided by eschewing a privacy notice altogether both because that would generally be competitive suicide and because that course would be deceptive in that it would entail a failure to disclose material facts ….
>
> In short, to the extent that privacy notices have been buried, incomplete, or otherwise ineffective—and they have been—the answer is to enhance efforts to enforce the "notice" model, not to replace it with a new framework.[9]

Another example of how the FTC's existing authority could be used more effectively bears emphasis. As Google noted in its Comments on the FTC Green paper, the "FTC has its own inquiry authority, even absent evidence of a violation, and its investigatory authority also serves what is effectively an audit function. As its track record demonstrates, the FTC utilizes its existing authority to ensure that companies are abiding by their fair information practice obligations and representations."[10] This is especially important given the emphasis placed by the Staff Report on the implementation of Privacy by Design and the use of Privacy Impact Assessments—both things which are susceptible to audits.

---

[7] Alexander Pope, *An Essay on Criticism*, 1709.

[8] Staff Report, *supra* note 2, at E-2-3, n. 4.

[9] *Id*., at E-1-2. *See also* First FTC Privacy Roundtable, Remarks of J. Howard Beales III, George Washington University, at 296-97.

[10] Google, *Comments of Google Inc.* 8, Jan. 28, 2011, http://www.ntia.doc.gov/comments/101214614-0614-01/attachments/FINALCommentsonDepartmentofCommercePrivacyGreenPaper%20%283%29.pdf (hereinafter "Google Comments").

Put simply, no one is arguing that the FTC should do nothing. But the agency should use its existing authority to the maximum extent possible before demanding new authority. There is good reason for caution about expanding the FTC's powers: The FTC is already unique in the vastness of its jurisdiction (over nearly the entire economy) and the flexibility of its powers (to punish "unfair" and "deceptive" trade practices). If untethered from the specific meanings of these terms, and certain procedural safeguards, the agency could essentially become a "second national legislature."[11] Giving the agency vast new powers over the use of data would, as more and more of our economy and society becomes dependent on the collection and use of data, risk repeating the agency's calamitous over-reach in the 1970s: The FTC so thoroughly abused its uniquely vast jurisdiction through an expansive conception of "unfairness" by, among other things, trying to ban advertising to children, that it was dubbed the "National Nanny" by the Washington Post, hardly a Thatcherite bastion.[12]

This is why, while the FTC may plan a valuable role as a partner in facilitating further improvement of privacy practices and technological empowerment of users, the agency should not attempt to play the lead role—as Google's comments on the Department of Commerce's Green paper explain:

> [T]he Department (including through a newly created Privacy Policy Office)—alone or in conjunction with relevant enforcement agencies such as the FTC—can convene working groups and synthesize recommendations to provide clear guidance on industry-specific measures needed to protect consumer privacy in a particular context or industry, and to update those recommendations as technology evolves. ***For this approach to be effective, however, the regulators must participate as an open-minded convener without preconceived assumptions as to the best outcome; otherwise, the process is merely government-driven regulation by another name***.[13]

II. The Role of FIPPS in a Dynamic World

To many in the privacy community, the Fair Information Practice Principles (FIPPS) are the "gold standard of privacy" and any suggestion that they not be enshrined in law to govern all aspects of privacy is nothing short of heresy. While no one would deny their value as a framework by which to conceptualize how to protect privacy, they do not answer the more important and difficult questions of how to reconcile privacy with other competing values in any and various situations facing those who must actually design, implement, evaluate and iterate privacy practices in the real world. Much like religious texts, the FIPPS can have great value, but are also too easily subject to overly orthodox, uncritical application by a priesthood of "true believers"—advocates who genuinely care about privacy and have the noblest of intentions

---

[11] Berin Szoka, The Progress & Freedom Foundation, *How Financial Overhaul Could Put the FTC on Steroids & Transform Internet Regulation Overnight*, Progress Snapshot 6.7, Mar. 2010, http://www.pff.org/issues-pubs/ps/2010/pdf/ps6.7-FTC_on_steroids.pdf.

[12] Editorial, *The FTC as National Nanny*, WASHINGTON POST, Mar. 1, 1978 at A14.

[13] Google Comments, *supra* note 10, at 9 (emphasis added).

about protecting others as consumers and citizens, but who downplay, or ignore, the difficulties of applying their doctrine in a world of competing values.

A. The FIPPS Must Be Applied Contextually, not Literally

Properly understanding the FIPPS requires understanding the assumptions on which they rest—just as studying any text requires an appreciation of its origins and how those translate to the present instance. Limiting access by government to particularly sensitive data (about health) raises a set of concerns for which the demands of FIPPS may well be appropriate. But the Internet is a far cry from the government-dominated healthcare sector of the 1970s. In contrast to the static world of "purpose specification," where "data minimization" means reducing possible harms at little cost, the dynamic world of the Internet is one where the most beneficial uses of data cannot be specified *ab initio* and where the minimization of data collection—the "precautionary principle" approach to privacy—comes at a significant cost.

Thus, for the FIPPS to be useful, they "must be appropriately tailored and relevant for their intended use,"[14] as Google argues—in other words, adapted to reflect the competing values at stake. The Interactive Advertising Bureau offers a simple illustration of the need for such adaptation, depending on the costs and harms at issue:

> [FIPPS'] data quality and integrity requirements are unnecessary in online advertising. The costs associated with building the infrastructure to permit access and correction rights for advertising and marketing data would significantly outweigh the supposed benefits from these rights. Inaccurate advertising and marketing data would at worse result in a less relevant advertising.[15]

B. Application of the FIPPS Must Allow for Ongoing Evolution

Google's comments on the Green Paper detail several outstanding examples of data collected for one purpose that were later used to develop services now used widely and without serious privacy concerns:

> Creative, even serendipitous re-use of collected data has enabled enormous advances in online products and services that enable creativity, education, the creation of businesses, and deeper social and political engagement. In Google's experience alone, purpose-compatible re-use of existing data has delivered enormous value to Google users and led to product improvements such as Gmail's priority inbox, automated spell checking, auto-complete, spam, fraud and virus protection tools, and the development of new services such as FluTrends and Translate. Mechanistic or overly prescriptive purpose specifications, data minimization and collection limitations, or use limitations

---

[14] *Id*. at 6.

[15] Interactive Advertising Bureau, *Letter RE: IAB's Comments* 7, Jan. 28, 2011, http://www.ntia.doc.gov/comments/101214614-0614-01/attachments/ACF2DA.pdf.

> would frustrate such economically and socially valuable innovation without protecting consumers from harm.[16]

As Facebook explains in its Green Paper comments, expectations themselves evolve alongside the technologies we use:

> As technology advances, individuals understand that their data may be used or made available in new ways. In the digital world in particular, users have come to understand and even expect that services will evolve and that companies will offer innovative new features that improve the online experience. The Department of Commerce's report, recognizing that creative reuses of existing information can lead to innovation but also cautioning that such innovative reuses should not come at the expense of user privacy, recommends a nuanced approach to the issue—one that weighs the benefits of the particular reuse against the harms and calibrates notice and consent requirements accordingly. Facebook believes that such an approach is necessary in lift of the many examples of reuse that have provided immense benefits to the public while producing little if any discernible harm. [17]

Because the beneficial uses of data co-evolve with privacy expectations, government must be careful not to foreclose innovation by attempting to freeze the status quo—such as by requiring companies to notify users, or receive consent, before ever using data in a new ways. As Facebook notes:

> While transparency is important, it must be implemented with due regard for the rapidly changing nature of online services and the realization that overly restrictive obligations hinder innovation. For example, the FTC recommends that companies obtain affirmative consent from users before using previously collected data in a "materially different manner" than described in an earlier privacy notice. While Facebook agrees that notice and consent may be appropriate for certain changes in data practices, it is essential to avoid interpreting the term "material" too restrictively. A restrictive interpretation could prevent companies from launching new features out of an uncertainty about whether those features would use data in a "materially different manner." Such an interpretation might have prevented features like the caller ID displays and Netflix recommendations described above from ever having been offered—a result that could hurt the future of the digital economy.[18]

---

[16] Google Comments, *supra* note 10, at 7-8

[17] Facebook, Inc., *Letter Re: Commercial Data Privacy and Innovation in the Internet Economy* 7, Jan. 28, 2011, http://www.ntia.doc.gov/comments/101214614-0614-01/attachments/FINALCommentsonDepartmentofCommercePrivacyGreenPaper%20%283%29.pdf (hereinafter "Facebook Comments").

[18] *Id*. at 9.

C. The Danger of Regulatory Capture

Attempting to impose a rigid regulatory regime grounded in FIPPS on a dynamic world is particularly dangerous because of the rapid, and accelerating, pace of technological change online. The problem is not simply that government will struggle to keep pace (it certainly will), but that policymakers must necessarily rely on the companies they regulate to understand the basic facts of new technologies and the (privacy) issues they create. As Tim Wu explains in *The Master Switch*, this reliance by regulator on regulatee means that the latter will inevitably attempt to capture regulation by casting narratives that skew to their advantage:

> The government can act only on the basis of what it understands to be established fact. Much of what is called lobbying must actually be recognized as a campaign to establish, as conventional wisdom, the "right" facts, whether pertaining to climate change, the advantages of charter schools, or the ideal technology for broadcasting. Much of the work of Washington lobbyists is simply an effort to control the conversation surrounding an issue, and new technologies are no exception.[19]

It was through such "fact-establishing" that, as Wu explains, the established incumbents of AM radio used the FCC as a weapon against competition by technologically superior FM radio in the 1930s and 40s.

Privacy regulation is no different from any other form of regulation, and is just as likely to be captured by special interests. Commissioner Rosch notes that such regulatory capture will occur not just when the FTC attempts to regulate outright, but also when it attempts to drive self-regulation:

> the self-regulation that is championed in this area …. may constitute a way for a powerful, well-entrenched competitor to raise the bar so as to create an entry barrier to a rival that may constrain the exercise of undue power[20]

The reality of regulatory capture is yet another reason for exercising caution in both regulating and attempting to shape self-regulation.

III. "Do Not Track"

Last week, Rep. Jackie Speier introduced legislation that would require the FTC to establish standards for a "Do Not Track" mechanism and require online data collectors to obey consumer opt-outs through such a tool.[21] In principle, a "Do Not Track" mechanism could enhance consumer empowerment, giving users the capacity to choose for themselves whether they want behavioral advertising. Such user empowerment tools are superior to restrictive defaults

---

[19] TIM WU, THE MASTER SWITCH, at 130 (201).

[20] Staff Report, *supra* note 2, at E-3.

[21] Congresswoman Jackie Speier, Do Not Track Me Online Act, Feb. 2, 2011, http://speier.house.gov/uploads/Do%20Not%20Track%20Me%20Online%20Act.pdf.

based on the paternalistic assumption that users won't make the "right" choice, no matter how easy that choice is to make. (Of course privacy zealots believe the "right" choice is always to minimize the collection and use of data, because data is dangerous.) But, as with so many things, the devil lies in the details. Even supporters of a "Do Not Track" mechanism should recognize that it would be premature for any technological mandate in this area, for three reasons:

First, markets *are* working. In the past, regulatory advocates insisted government must intervene immediately because, they argued, markets had failed to address privacy concerns. But just days before Rep. Speier introduced her legislation, Microsoft and the Mozilla Foundation launched "do-not-track" tools in new versions of their Internet browsers: Internet Explorer 9[22] and Firefox 4,[23] while Google launched a tool as an add-on for Chrome.[24]

Second, the FTC already has the authority to enforce promises made by data collectors to comply with the wishes of users who express a preference not to be tracked via a "Do Not Track" mechanism. Regulatory advocates, of course, will argue that too few companies will make such promises for this marketplace response to be effective and, therefore, that government must not only enforce such promises, but also mandate compliance with users' "Do Not Track" preferences—and also perhaps mandate use of a "Do Not Track" standard by browser-makers. But it is simply too soon to say how this will develop. And even if it does turn out that many data collectors remain silent on honoring "Do Not Track," other technologies such as Microsoft's variant may simply allow users to block *all* content from such data collectors—including tracking code.

In any event, the technical details of a "Do Not Track" mechanism must be allowed to evolve over time. We cannot expect a workable "Do Not Track" mechanism to simply spring into being overnight—much as people imagined, for centuries after Aristotle, that life was capable of "spontaneous generation." Instead, Ultimately, it is the Internet's existing standards-setting bodies (*e.g.*, W3C, IETF), not Congress or the FTC, that have the expertise to resolve such differences and make a "Do Not Track" mechanism work for both consumers and publishers, as well as advertisers and ad networks. Specifically, that will require some degree of standardization of the following, among other things:

- The definition of "tracking";
- The interface by which users activate and configure the "Do Not Track" mechanism; and
- The process by which websites respond to the mechanism and negotiate with users who want to opt-out of tracking for access to content.

---

[22] Sean Hollister, Internet Explorer 9 RC Now Available to Download, Tracking Protection in Tow (Update), Feb. 10, 2011, http://www.engadget.com/2011/02/10/internet-explorer-9-rc-now-available-to-download-tracking-prote/.

[23] *Mozilla Firefox 4 Beta, Now Including "Do Not Track" Capabilities*, The Mozilla Blog, Feb. 8, 2011, http://blog.mozilla.com/blog/2011/02/08/mozilla-firefox-4-beta-now-including-do-not-track-capabilities/.

[24] *See* https://chrome.google.com/webstore/detail/hhnjdplhmcnkiecampfdgfjilccfpfoe.

A. Possible Economic Consequences

Jonathan Mayer, of Stanford's Center for Internet & Society, insists that we need not fret about the economic consequences of "Do Not Track" because, among other reasons, behavioral advertising revenue is a relatively small share of total U.S. online advertising spending: just 4%, he insists.[25] But his comparison mixes apples and oranges: The relevant comparison is not behavioral advertising not to total online advertising revenue (including search advertising spending), but to spending on *display* advertising (advertising sold by websites next to their content): Behavioral advertising spending in 2010 represented roughly 20% of total display ad spending and that ratio is expected to grow.[26]

Regardless, as Ben Kunz explains, the question is not merely how much revenue is available for ad-supported media, but how that revenue is distributed:

> Like the publications of the past century, a given website has always been a proxy for an audience target. Alas for the big publishers, good data on audiences has meant that smart marketers could leave big, expensive sites behind. So in perhaps the biggest revolution of Internet marketing, the more data you can collect about today's customers, the cheaper online advertising gets ….
>
> If the FTC pushes Do Not Track through Congress, it will send billions to The Wall Street Journal (NWS), Forbes.com, iVillage.com, and even Bloomberg Businessweek because marketers will be forced to put ad dollars on those sites. In the absence of data, advertisers will have to make assumptions about who reads content. The top content will win.[27]

In other words, adoption of a "Do Not Track" mechanism could have significant consequences for the structure of the media sector. Some, like *Cult of the Amateur* author Andrew Keen, might argue that this redirection of revenue towards larger, better established websites (and offline to traditional media) is desirable to preserve elite, "professional" media. Others would counter, without (necessarily) denying the value of traditional media, that the "Long Tail" of websites disadvantaged by "Do Not Track" represent diversity, creativity and the "laboratories" of media's future (think Huffington Post). The important point is not which side has the better argument, but that such arguments—over picking winners and losers—are the very hallmark of industrial planning.

---

[25] Jonathan Mayer, *Do Not Track Is No Threat to Ad-Supported Businesses*, Jan. 20, 2011, http://cyberlaw.stanford.edu/node/6592.

[26] Network Advertising Initiative, Study Finds Behaviorally-Targeted Ads More Than Twice As Valuable, Twice As Effective As Non-Targeted Online Ads, http://www.networkadvertising.org/pdfs/NAI_Beales_Release.pdf ("Behaviorally-targeted ads accounted for 17.9% of respondents' advertising revenue, with revenue increasing from 16.2% in Q1 to 19.4% in Q4 2009.").

[27] Ben Kunz, *The $8 Billion Do Not Track Prize*, BLOOMBURG BUSINESSWEEK, Dec. 22, 2010, http://www.businessweek.com/technology/content/dec2010/tc20101222_392883.htm.

Mayer insists "[a]d-supported businesses could ask—or possibly require—Do Not Track users to allow third-party behavioral advertising."[28] In other words, he argues that the market will solve this problem—to be precise, the "market for privacy" created by empowering users to forbid websites to use their data for behavioral advertising purposes. Perhaps. But how will that work? And how well?

Again, I am deeply sympathetic to the concept of creating a privacy marketplace. Adam Thierer and I have, from the start of our work, argued for recognition of the value exchange underlying the Internet ecosystem: Publishers offer free content and services and in exchange, users offer a share of their attention (viewing those ads) as well as information about where their attention is likely to go (making those ads more relevant).

Yet we simply do not know how this new marketplace will evolve as today's implicit *quid pro quo* becomes, or is forced to become, explicit. Thus, government must be cautious when it attempts to design that marketplace from the top down through regulation (as would happen under the bill introduced last week by Rep. Jackie Speier). The same is true when government acts more subtly, using the bully pulpit to intimidate industry (as Chairman Leibowitz has essentially done since calling for "Do Not Track" in Congressional testimony last July[29]). Much as I enjoy the rich irony of seeing those who are rarely thought of as free-marketeers essentially asserting that "markets" will simply, and quickly, "figure it out," I am less sanguine. The hallmark of a true free-marketeer is not a belief that markets work perfectly; indeed, it is precisely the opposite: an understanding that "failure" occurs all the time, but that government failure is generally worse, in terms of its full consequences, than "market" failure.

The first part of that lesson comes especially from the work of the economist Ronald Coase, who did more to teach us "how little [we] really know about what [we] imagine [we] can design" than perhaps anyone. Coase won his Nobel Prize for explaining that the way property right are allocated and markets are structured determines the outcome of marketplace transactions. For example, a rule that farmers bear the cost of stopping rancher's cattle from grazing on their farms by constructing fences will produce different outcomes—not merely different allocations of costs—from the opposite rule.

Coase's key insight was that, in a perfectly efficient market, the outcome would not depend upon such rules: To put this in terms of the privacy debate, the choice between, say, an opt-out rule and an opt-in rule for the collection or use of a particular kind of data (essentially a property right) *would have no consequence* because the parties to the transaction (say, website users and website owners) would express their "true" preferences perfectly, effortlessly and costlessly. But, of course, such frictionless nirvanas do not exist. The real world is defined by what Coase called "transactions costs": search and information costs, bargaining and decision costs, policing and enforcement costs.

---

[28] Mayer, *supra* note 25.

[29] Juliana Gruenwald, *FTC Weighs 'Do Not Track' List*, NationalJournal, July 27, 2010, http://techdailydose.nationaljournal.com/2010/07/ftc-weighs-do-not-track-list.php.

The transaction costs of implementing a "Do Not Track" mechanism as something other than pure free-riding (no-cost opt-outs being, ultimately, unsustainable) are considerable: someone must design interfaces that make it clear to the user what their choice means, the user must consume that information and make a choice about tracking, websites must decide how to respond to various possible choices and be able to respond to users in various ways through an interface that is intelligible to users, and so on—all for what might seem like a "simple" negotiation to take place.

These problems are certainly not insurmountable—and, again, with the right engineering and thoughtful user interface design a "Do Not Track" mechanism could well prove a useful tool for expressing user choice. But when we look at the world through Coase's eyes, we begin to understand that how mechanism design can radically can outcomes (in this case, funding for websites. Indeed, the costs of building and operating a market for privacy—measured in time as well as money—could well swamp the value produced by that market. Clearly, website publishers currently write-off ad-blocking as an acceptable loss because it would cost them more to fix the problem (*e.g.*, by charging users who block ads) than they would gain in revenue by doing so. The question is: how high is that threshold? And how much total revenue will be lost even when publishers are able to get *some* users to pay *something*? These are just some of the questions that must be answered before government inserts itself into the evolution of user choice mechanisms.

B. More Sophisticated "Do Not Track" Mechanisms

David Vladeck, director of the FTC's Bureau of Consumer Protection, recently made clear in Congressional testimony that the agency ultimately wants a much more granular choice mechanism:

> We therefore urge Congress to consider whether a uniform and comprehensive choice mechanism should include an option that enables consumers to control the types of advertising they want to receive and the types of data they are willing to have collected about them, in addition to providing the option to opt out completely.[30]

In many respects, this is admirable. One of the significant drawbacks to the "Do Not Track" mechanism as implemented in Firefox 4 is that it allows only the expression of a single preference not to be "tracked" across the board—as this screen capture illustrates:

[30] David Vladeck, *Prepared Statement on Do Not Track Before the Committee on Energy and Commerce, Subcommittee on Commerce, Trade, and Consumer Protection, United States House of Representatives*, Dec. 2, 2010, http://www.ftc.gov/opa/2010/12/dnttestimony.shtm.

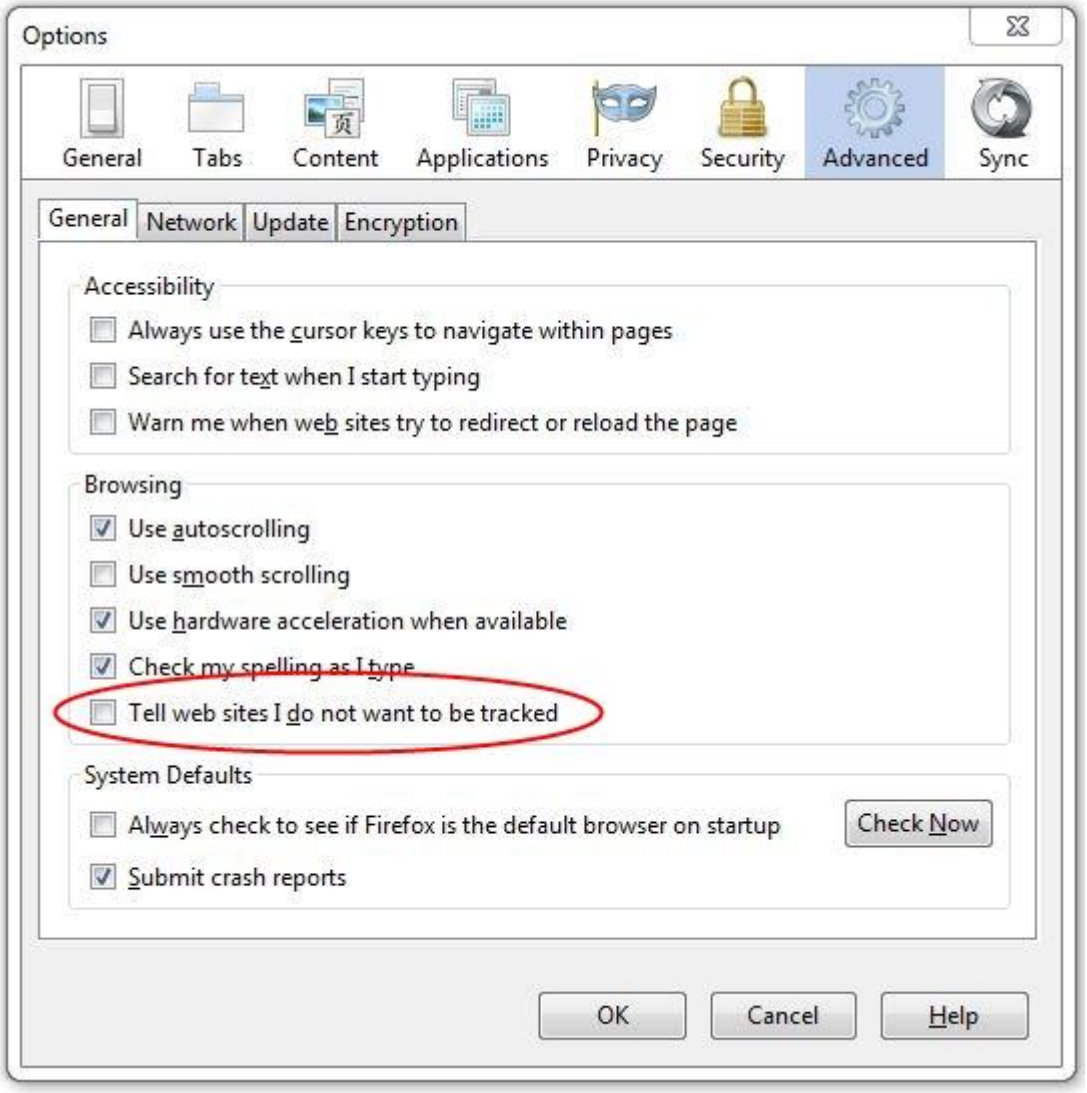

The Electronic Foundation's Green paper comments laud the promise of "Do Not Track" as the first of a potential new generation of user empowerment tools that could give effect to a core FIPPS principle better than a simple legal mandate:

> DNT is just one example of the way that technical measures may improve purpose-related disclosure. DNT is a consumer-expressed preference that says the user's browser information may be used for sending content to the user, but not for recording the user's reading habits. Over time, we believe that similar standards should and will be developed for other kinds of purpose specification.[31]

[31] Electronic Frontier Foundation, Comments to the Department of Commerce Internet Policy Task Force 5-6, http://www.ntia.doc.gov/comments/101214614-0614-01/attachments/ACF2D4.pdf.

But this vision of achieving a core concept of FIPPS—purpose specification—by user empowerment raises significant practical questions. How would more complicated mechanisms actually work? How would websites respond to a wider array of expressions of user intent about what may be done with their data? What would the resulting "marketplace for privacy" look like? What would be the transactions costs of implementing such a marketplace? Would the value generated in such a marketplace sufficiently outweigh transactions costs that the marketplace could continue to sustain ad-supported content and services?

Most pointedly: Why should we believe that the FTC is best suited to answering these difficult questions? I can only hope that the agency will not attempt to answer these questions on its own, but instead rely on the marketplace to develop clearer answers. I will readily join hands with EFF in celebrating user empowerment tools in the privacy context (just as we do in the context of online child protection), but I remain skeptical about the wisdom of having government design such tools, concerned about how well such tools will work, and what their costs will be.

C. Microsoft's "Do Not Load" Mechanism

In the end, Microsoft's IE9 mechanism—which might more accurately be dubbed "Do Not Load,"—might well moot the debate over what tracking means by empowering users to block any content that loads tracking elements whose data collection, use, access or security practices are deemed inadequate by the maker of the Tracking Protection List (TPL) installed by the user.

In principle, such a mechanism is highly compelling for two reasons: First, it is self-enforcing, because the browser simply does not load blacklisted content, rather than relying on a third party to respect a preference expressed in a heading, someone else to detect violations of that preference, an effective punishment, *etc*. As noted above, such a mechanism could someday work in conjunction with a "Do Not Track" mechanism such as that offered by Firefox by blocking content from companies that do not commit to respecting the "Do Not Track" suspenders—a promise the FTC, in turn, would enforce.

Second, the way Microsoft has designed their mechanism is directly analogous to parental control tools that empower the parent to implement their preferences by subscribing to the white list or black list of a trusted third party. In principle, such subscription tools can empower us to make effective tools about complex problems by outsourcing the decisions to trusted third parties—be they large or small, for-profit or non-profit, from corporations to churches to privacy advocacy groups.

Yet "Do Not Load" also raises significant questions. Again, how would a marketplace for privacy actually work to empower publishers to condition access to their content? What would be the costs of building such a marketplace?

And how could such a mechanism be used to manipulate the online market for content and services? After all, "Do Not Load" is simply a powerful tool for blocking content that, in the hands of third parties whose interests do not fully align with users, could be used for great mischief. Microsoft has, wisely, abstained from writing its own TPLs—instead choosing simply

to build the mechanism and let others write TPLs. But what if a TPL were used to block a competitor's content? For example, suppose a computer OEM pre-installed certain TPLs on a browser shipped to the consumer that simply removed page elements served by competitors. This could, in theory, cause Facebook's "Like" button to simply disappear from all websites, for example.

The purpose of this hypothetical is not to trot out a "parade of horribles" that will necessarily follow any "Do Not Track" effort, but to illustrate how little we understand about the real-world consequences of such user mechanisms, and to highlight how dependent those consequences are on mechanism design. Much imaginable market manipulation could be largely addressed by designing an architecture that is transparent to the user. Here are just a few of the questions that ought to be asked about a "Do Not Load" mechanism:

- Could TPLs come pre-installed?
- Would users see the contents of the lists?
- Will users know if/when lists have been updated?
- How will users be informed about the contents of a TPL white list they might be asked to install by a website that is attempting to negotiate with them over access to content?

With so many questions about two radically different user choice mechanisms, it would be premature for the FTC to even begin to contemplate technological mandates in this area—as Rep. Speier's proposed legislation would *require* the agency to do.

D. "Tracking Neutrality"

A very different sort of "neutrality" concern has been raised—over how publishers interact with users. The discussion above concerns how, as a practical matter, websites would respond to privacy-sensitive users who opted out of tracking through the "Do Not Track" header and what the consequences of that back-and-forth might be. A true "privacy marketplace" would be based on empowering users to implement their privacy preferences in a meaningful way, while also empowering website publishers to respond to opt-outs as they see fit.

But this, of course, presumes that websites would be free to condition access to their content on receiving permission to "track" users for advertising purposes or, failing that, charge for their content, or otherwise discourage users from opting out (such as by showing them more ads, or "interstitial" ads with a count-down before they can access a desired page). Harlan Yu, a researcher at Princeton's Center for Information & Society, would allow websites to do so—but only so long as their "discrimination" against privacy-sensitive users was "reasonable":

> nothing would prevent sites from offering limited content or features to users who choose to opt-out of tracking. One could imagine a divided Web, where a user who turns on the x-notrack ["Do Not Track"] header for all HTTP connections—i.e. a blanket opt-out—would essentially turn off many of the useful features on the Web.
>
> By being more judicious in the use of x-notrack, a user could permit silos of first-party tracking in exchange for individual feature-rich sites, while limiting

> widespread tracking by third parties. But many third parties offer useful services, like embedding videos or integrating social media features, and they might require that users disable x-notrack in order to access their services. Users could theoretically make a privacy choice for each third party, but such a reality seems antithetical to the motivations behind Do Not Track: to give consumers an easy mechanism to opt-out of harmful online tracking in one fell swoop.
>
> ***The FTC could potentially remedy this scenario by including some provision for "tracking neutrality," which would prohibit sites from unnecessarily discriminating against a user's choice not to be tracked.*** I won't get into the details here, but suffice it to say that crafting a narrow yet effective neutrality provision would be highly contentious.[32]

Yu likely won't be the only one to suggest such restrictions, which will likely find support from those in the "free culture" movement who generally do not accept that those who produce content, or offer a service, have every right (subject to fair use, consumer deception laws and antitrust) to condition or restrict access to that content/service. The Staff Report itself leaves the door open to such proposals, as Commission Rosch notes—and rightly rejects:

> insofar as the Report could be read as suggesting a ban on "take it or leave it" options (see Report at 60), again, clear and conspicuous disclosure is the most appropriate way to deal with such an option. I question whether such a ban would be constitutional and am also concerned about the impact of a ban on innovation.[33]

This serves merely to illustrate one dimension of the "And then what?" approach policymakers should follow in understanding the many and various consequences of pushing a "Do Not Track" mechanism. Harlan is clearly right about one thing: "Do Not Track," as the title of his blog post says, is "Not as Simple as it Sounds."

E. Metrics for Success

In the end, perhaps the most important question to be asked about "Do Not Track" is: What are the metrics for success? When many "Do Not Track" advocates draw analogies to the "Do Not Call" registry, they imply that success would look similar in both cases: adoption by a majority for users. But, returning to the alternative framework for approaching privacy outlined above, we cannot know what value users really place on privacy until we see their preferences revealed in the marketplace when they must choose from among competing variables. In other words, we really do not know how many people would choose to enact "Do Not Track" when presented with a choice among clear alternatives: allow tracking, move on to another site, or pay some cost in terms of additional advertising, or payment for content. So, how will we know

---

[32] Harlan Yu, *Do Not Track: Not as Simple as it Sounds*, Freedom to Tinker, Aug. 10, 2010, http://www.freedom-to-tinker.com/blog/harlanyu/do-not-track-not-simple-it-sounds (emphasis added).

[33] Staff Report, *supra* note 2, at E-6.

how much adoption is enough? Or will it be the availability of useful tools that matters, regardless of how many people use them? And how will we measure the costs of such mechanism? The FTC has proposed an admirable standard:

> [A]ny such [Do Not Track] mechanism should not undermine the benefits that online behavioral advertising has to offer, by funding online content and services and providing personalized advertisements that many consumers value.[34]

What kind of empirical evidence would actually satisfy us that such a standard has been met?

IV. "~~Elvis~~ *The First Amendment* Has Left the Building!"

Whatever government does in regulating the use and collection of data online cannot be done without regard to the First Amendment, because (i) online "privacy" regulation is the regulation of how data flows in the Internet ecosystem, (ii) those data flows are essential to the tailoring, delivery and funding of online speech, and thus (iii) restrictions on the flow of data are, to varying degrees and each in their own ways, restrictions on speech itself. This is not to say that government may do nothing, but that we must understand how any particular proposed government intervention affects online speech, decide what level of First Amendment scrutiny applies, and then ask whether the government has met its burden to satisfy that scrutiny.

Sadly, the First Amendment seems to be almost entirely absent from the general drive towards increased privacy regulation. Nowhere does the FTC staff report mention "free speech" or the "First Amendment."[35] Only Commissioner Rosch mentions concerns about the First Amendment, noting that it might be unconstitutional for the FTC to ban "'take it or leave it' options" by which website publishers would refuse to make their content available unless users accepted tracking.[36] Yes, indeed, dictating to publishers on what terms they may make their content available would be the ad-supported ("free") content world's equivalent of price controls. Turning media providers into public utilities that must provide content as "common carriers" to all visitors, regardless of whether those visitors contribute to the business model that funds free content, would obviously impinge on the First Amendment rights of publishers.

But this is only the most extreme example of a more general First Amendment problem raised by privacy regulations: When government regulates the use of data for advertising purposes, it necessarily affects the funding available to ad-supported publishers, as noted above.

---

[34] Staff Report, *supra* note 2, at 67.

[35] Even the ACLU, perhaps America's most stalwart defender of free speech, seem oblivious to the integral relationship between free speech and the flow of information: The FTC cites their primer *Privacy and Free Speech: It's Good for Business*, which, despite its name is *not* about the relationship *between* privacy and free speech, but about how companies can suffer in the marketplace by invading privacy or interfering with free speech. Staff Report at 45 (citing ACLU, *Privacy and Free Speech: It's Good for Business*, http://www.aclunc.org/docs/technology/privacy_and_free_speech_it's_good_for_business.pdf). While this is, indeed, a core argument that market forces will drive companies to self-regulate, it misses the larger connection between free speech and privacy.

[36] Staff Report, *supra* note 2, at E-6.

More generally, speech and privacy are but two sides of the same coin. After all, what is your "right to privacy" but a right to stop me from observing you and speaking about you?[37] Thus, when government restricts the collection of information, it also restricts the processing and reporting of information—also known as "reporting." This point is commonly recognized, yet few people think through the implications of data regulations for online speech. The simple truth is that online speech is only as effective as it is "targeted" to a particular audience—and that effective "targeting" requires useful data about the likely interests of a potential reader/listener/viewer. This is as true for companies that buy online ads for toothpaste as it is for political candidates and non-profit causes that attempt to reach voters, supporters, donors and volunteers through online media.

If government limits the ability to speak effectively online, whether through direct regulation or indirect pressure, it necessarily implicates the First Amendment. The difficulty facing the FTC in this area lies in the nature of online speech platforms: Past laws regulating the Internet (*e.g.*, COPPA, COPA) have attempted to avoid First Amendment problems by exempting non-commercial websites (and thus avoiding the strict scrutiny standard), but this approach breaks down in a world where online speech flows not from individual websites, but through *platforms*. For instance, if government regulation reduces the data available to target an ad through an ad network or a message through a social network, that regulation necessarily falls on both commercial speakers (the toothpaste ad) and non-commercial speakers (the political or message ad). There is no easy way to "carve out" more highly protected non-commercial speech, because data regulations burden the platforms that carry *all* online speech.

V. An Alternative Framework for Approaching Privacy

So, how *should* policymakers and companies approach privacy, in deciding how to apply FIPPS and other ideas about privacy in the real world? As I argued in my earlier filing on the FTC's Privacy Roundtables,[38] any discussion about regulating the collection, sharing, and use of consumer information online must begin by recognizing the following:

- Privacy is "the subjective condition that people experience when they have power to control information about themselves and when they exercise that power consistent with their interests and values."[39]
- As such, privacy is not a monolith but varies from user to user, from application to application and situation to situation.

---

[37] Eugene Volokh, *Freedom of Speech and Information Privacy: The Troubling Implications of a Right to Stop People From Speaking About You*, 52 STANFORD L. REV. 1049 (2000), http://www.pff.org/issuespubs/pops/pop7.15freedomofspeech.pdf.

[38] Berin Szoka, *Privacy Trade-Offs: How Further Regulation Could Diminish Consumer Choice, Raise Prices, Quash Digital Innovation & Curtail Free Speech,* Comments to the FTC Privacy Roundtables (Dec. 7, 2009), Comment, Project No. P095416, http://www.pff.org/issues-pubs/filings/2009/111009-FTC-privacy-workshop-filing.pdf.

[39] "Properly defined, privacy is the subjective condition people experience when they have power to control information about themselves." Jim Harper, Cato Institute, *Understanding Privacy–and the Real Threats to It*, Cato Institute Policy Analysis No. 520, Aug. 4, 2004, http://www.cato.org/pub_display.php?pub_id=1652.

- *There is no free lunch*: We cannot escape the trade-off between locking down information and the many benefits for consumers of the free flow of information.
- In particular, tailored advertising offers significant benefits to users, including potentially enormous increases in funding for the publishers of ad-supported content and services, improved information about products in general, and lower prices and increased innovation throughout the economy.
- Tailored advertising increases the effectiveness of speech of all kinds, whether the advertiser is "selling" products, services, ideas, political candidates or communities.

With these considerations in mind, policymakers should always look for the "least restrictive" means available to address clear harms—in the broad, but still provable, sense Commissioner Rosch talks about harm. Beyond preventing unfair and deceptive trade practices by the companies that use and collect online data, government can also play a vital role in protecting consumers from real harms that flow from the use of their data, such as the use of personal data to make decisions about credit. Government may even play a proper role in supporting education about privacy risks and promoting technical tools that empower consumers to make more effective decisions about their own privacy—just as it has done with parental empowerment solutions to address concerns about online child safety and protection.

But as in that context, where the courts insist on such a "least-restrictive means" test as a matter of First Amendment doctrine, we have argued consistently for the following layered approach to concerns about online privacy.[40] Government should:

1. ***Erect*** a higher "Wall of Separation between Web and State" by increasing Americans' protection from government access to their personal data—thus bringing the Fourth Amendment into the Digital Age.
2. ***Educate*** users about privacy risks and data management in general as well as specific practices and policies for safer computing.
3. ***Empower*** users to implement their preferences about the real-world trade-offs between privacy and other values as easily as possible.
4. ***Enhance*** self-regulation by industry sectors and companies to integrate with user education and empowerment tools (*e.g.*, respecting evolving consumer choice mechanisms).
5. ***Enforce*** existing laws against unfair and deceptive trade practices as well as state privacy tort laws.

[40] *See, e.g.,* Berin Szoka & Adam Thierer, *Online Advertising & User Privacy: Principles to Guide the Debate*, Progress Snapshot 4.19, Sept. 2008, available at http://www.pff.org/issues-pubs/ps/2008/ps4.19onlinetargeting.html.

VI. Conclusion

The approach I propose above might be called "conservative." In one sense, it is just the opposite: an argument that we must embrace change in a dynamic world, and that we cannot maintain the technological status quo.[41]

But in another sense, this approach does harken back to Edmund Burke's "conservatism" of prudence. Burke, in general, argued against the absolutist radicalism of the French Revolution's Jacobin elements, earning him the caricature as a purely reactionary champion of the status quo and its established interests. Yet Burke was, in fact, a great champion of reform in his day—and the leading defender of the American colonists' grievances against British oppression before the Revolution. "A State without the means of some change is without the means of its conservation," Burke wrote.[42]

The same is true of the Internet ecosystem when it comes to improving data collection practices and user empowerment: Ultimately, we *do* need better user empowerment tools like "Do Not Track." Yet there is a middle ground between doing nothing and the insistence of those privacy Jacobins who demand immediate, sweeping intervention, no matter its costs, because privacy is a "fundamental right" that must be protected at any cost—"*Fiat justitia ruat coelom*," as Latin-loving lawyers say: "May justice be done though the heavens fall."

The FTC has a key role to play in this process, as Commissioner Rosch has argued. Yet "Rome was not built in a day," and neither will be a sustainable privacy marketplace that works for both consumers and publishers, as well as the advertisers and ad networks who "keep the party going" for everyone. Where the market process of discovery through innovation is working, government should not interfere. That process is well underway with "Do Not Track." There is much to lose by rushing forward. The FTC should follow Burke's maxim: "Our patience will achieve more than our force."[43] Participating in the inter-agency working process outlined by the Department of Commerce in its Green Paper will be a good way for the FTC to put its "patience" to good use—helping to improve best practices, but not dictating them.

---

41 *See generally* VIRGINIA POSTREL, THE FUTURE AND ITS ENEMIES (1998).

42 Edmund Burke, Reflections on the Revolution in France (1790), available at http://www.constitution.org/eb/rev_fran.htm.

43 *Id.*