Unjustified By Design: Unfairness and the FTC's Regulation of Privacy and Data Security

Randal C. Picker[*]

Nature and regulators are very much alike: both abhor a vacuum. The shift of life online has brought with it an unusual power to accumulate information about consumers, though, with the emergence of new technologies such as facial recognition technology, the opportunities for information gathering in physical space are increasing as well. We should expect substantial changes in technology to give rise to conflicts over the appropriate boundaries of that technology and calls for regulation.

And regulators have indeed responded. In February, 2012, the White House issued its consumer data privacy report with a call for a new Consumer Privacy Bill of Rights.[1] The U.S. Federal Trade Commission has been particularly active issuing a major consumer privacy report in March, 2012 and subsequent reports in September, 2012 and October, 2012 offering mobile app guidelines and addressing the use of facial recognition technology.[2] And the FTC has done much more than just issue reports: the

 James Parker Hall Distinguished Service Professor of Law and Senior Fellow, The Computation Institute of the University of Chicago and Argonne National Laboratory. I thank the Sarah Scaife Foundation and the Paul Leffmann Fund for their generous research support. A prior version of this paper was given in December, 2012 at a conference at George Mason Law School under the title "Structuring Competition in Privacy" and I thank the participants in that conference for their comments. I have received an honorarium from the Law & Economics Center at George Mason University School of Law in connection with the preparation of this article.

[1] The White House, Consumer Data Privacy in a Networked World: A Framework for Protecting Privacy and Promoting Innovation in the Global Digital Economy, February, 2012 (online at xxx.)

[2] U.S. Federal Trade Commission, Protecting Consumer Privacy in an Era of Rapid Change, March, 2012 (online at xxx); U.S. Federal Trade Commission, FTC Publishes Guide to Help Mobile Developers Observe Truth-in-Advertising, Privacy Principles, September 5, 2012 (online at http://www.ftc.gov/opa/2012/09/mobileapps.shtm); U.S. Federal Trade Commission, Facing Facts: Best Practices for Common Uses of Facial Recognition Technologies, October, 2012 (online at xxx).

breadth of the underlying statutes under which the FTC operates and the natural mistakes that firms will make in a rapidly developing industry have made it possible for the FTC to move aggressively into direct privacy regulation.[3]

And it isn't just the U.S. federal government that is moving forward on privacy regulation. Without even considering developments in the EU—and with most of what is going on in privacy occurring on the Internet or through apps on devices like tablets and smartphones, we really need to focus on the world market—California has moved to extended its online privacy regime to apps.[4] At 1 million mobile apps on the iOS and Android platforms and a potential fine of $2,500 per non-complying download, California may have figured out how to solve its budget problems.[5]

In this paper, I focus on a core principle around much of which the FTC's recent analysis is built, namely that firms should organize their efforts in building products to promote privacy by design or, perhaps with parallel effect, to promote security by design. I think that there are good reasons to be skeptical of those principles, both on a freestanding basis and then as measured against the statutory standards that restrict FTC authority. The regulatory approach seized upon by the FTC—a mix of reports and consent decrees—has largely allowed the FTC to sidestep the

[3] Current FTC Chairwoman Edith Ramirez set forth an overview of the FTC's extensive privacy-related actions in her June, 2011 testimony before a House subcommittee. See Federal Trade Commission, Prepared Statement on Data Security, June 15, 2011 (online at http://www.ftc.gov/os/testimony/110615datasecurityhouse.pdf)

[4] See State of California Department of Justice, Office of the Attorney General, Attorney General Kamala D. Harris Secures Global Agreement to Strengthen Privacy Protections for Users of Mobile Applications, February 22, 2012 (online at http://oag.ca.gov/news/press-releases/attorney-general-kamala-d-harris-secures-global-agreement-strengthen-privacy).

[5] State of California Department of Justice, Office of the Attorney General, Mobile Applications and Mobile Privacy Fact Sheet (online at http://oag.ca.gov/system/files/attachments/press_releases/n2630_updated_mobile_apps_info.pdf); Joe Mullin, CA to app devs: get privacy policies or risk $2500-per-download fines, xxx.

statutory restrictions that it would otherwise face. The FTC hasn't faced a serious judicial inquiry into its claimed powers as to privacy and data security, though the pending *Wyndham* case may do that.

In Section I of the paper, I consider how the FTC is using its Section 5 authority to regulate privacy.[6] It is the open-ended nature of that section which makes it easy for the FTC to move aggressively in regulating new areas. Much of the FTC's direct regulatory efforts to date have piggybacked on the privacy disclosures required by other law or made voluntarily by firms themselves. This puts the FTC in the posture of engaging in purely after-the-fact, one-by-one regulation of firms and doesn't push the FTC to articulate broader standards. And many of the situations are resolved through settlements such that the underlying issues aren't tested through litigation. The FTC could issue substantive rules as it has done in the past in other areas, but as Congress has repeatedly amended the statutes to create a demanding standard for many of the rules that the FTC might issue, the FTC has moved to using the report process described above. The reports sidestep the statutory standards that the FTC would otherwise face and make it possible for the FTC to issue non-rule "rules," rules that the FTC hopes will shape the relevant industry but without obvious direct legal effect. Of course, the line between actual rules and faux rules may not be clear to all involved and the FTC may indeed welcome that ambiguity. The privacy-by-design and security-by-design ideas are good examples of these non-rule rules.

In Section II of the paper, I focus on what should we expect in a competitive market in consumer data. Privacy-attentive consumers will be presented with choices that they will attend to, either through direct competition through data limits or through

[6] Note that I am not addressing the other more specific authority that the FTC has to regulate activities related to privacy, such as its authority under the Fair Credit Reporting Act. See, e.g., Federal Trade Commission, FTC Warns Data Broker Operations of Possible Privacy Violations, May 7, 2013 (online at http://www.ftc.gov/opa/2013/05/databroker.shtm).

personalization signals to enable choices. Even privacy-insensitive consumers will benefit from competition as firms will value the data that those consumers will provide and will offer additional value to those consumers to attract them to their services. But we should expect firms to overconsume data as it were, meaning to capture data from privacy-inattentive consumers where the value to the firms of receiving that data is less than the value to the consumers giving up the data.

In Section III of the paper, I consider mechanisms for addressing the overcapture of data from privacy-inattentive consumers. Of course, those may just be consumers who don't value privacy very much, so I look for metrics to assess whether consumers are interacting with transparency tools—such as data collection icons and personalization signals—in the way that we might expect. I then turn to considering the tools available to the government to perturb how consumers interact with these privacy signals. The government could require online services and apps to disclose more information—think the online equivalent of the FTC's octane or home-insulation rules—but a less centralized approach would be for the government itself to build disclosure apps available for downloading.

In Section IV of the paper, I consider a core part of the three-part framework put forward by the FTC in its March, 2012 privacy report, namely the requirement of privacy by design. I also look at the related notion of security by design at work in the FTC's recent settlement with HTC over smartphone data security.[7] In framing, as I do in the title of the paper, what the FTC is doing as to privacy and data security under the unfairness prong of Section 5 as unjustified by design, I mean two related ideas. First, the FTC hasn't really had to justify its approach. By proceeding outside of a rulemaking process and through uncontested consent orders, the FTC has sidestepped many of the procedural steps that help to test new regulatory approaches. It is that sense in which I regard what

[7] In the Matter of HTC America Inc., FTC File No. 122 3049, Feb 22, 2013 (online at http://www.ftc.gov/os/caselist/1223049/index.shtm).

the FTC is doing as unjustified. But I also mean more specifically that as the FTC edges away from a reliance on a deceptive practice framing of Section 5 privacy violations to one that embraces unfair acts or practices, it hasn't fully confronted the demanding standards of proof set forth in Section 5(n) that must be met if the FTC is to condemn an act or practice as unfair.

I. Regulating Privacy under Section 5

There are many basic questions in regulating privacy and data collection and I will touch upon some of them at various points, but I also want to be clear on what I am pushing to the side and not considering. So there is a general question about the extent to which we should prefer centralized regulatory approaches implemented through agencies or political actors such as state attorneys general versus more decentralized, common-law like approaches implemented through courts and centered on notions like harm, tort, implied contract and other notions.[8]

I will sidestep that issue here and instead will focus on the approaches available to the FTC for regulating privacy. And my interest isn't statutes such as the Fair Credit Reporting Act or Gramm-Leach-Bliley which set forth more particularized privacy-related powers, but is instead the FTC's Section 5 authority. The FTC is first and foremost an adjudicatory agency, where it exercises its authority when confronted with a particular controversy and that is true for privacy as well. The pattern of regulation here is clear: firms are often required to make privacy disclosures and they may not fulfill the statements made in those disclosures. Sometimes that seems willful, but in many other cases, firms just make mistakes. The FTC can sweep all of that into its

[8] On the general question of regulation versus tort-like liability, see Steven Shavell, A Fundamental Enforcement Cost Advantage Of the Negligence Rule over Regulation (working paper, 2013). And for an example of the common-law approach in action in the context of a class action over a data breach, see Anderson v. Hannaford Brothers Co., 659 F.3d 151 (1st Cir. 2011) (finding that plaintiffs had sufficiently alleged compensable harm and possible theories of liability sounding in negligence and implied contract in connection with data breach).

deceptive practices framework. The FTC could act more broadly through its rulemaking authority, as it has in the past most notably in its set of sixteen trade regulation rules (covering subjects as diverse as cooling-off periods for home purchases and labeling requirements for home insulation).[9]

But rather than enter into more general rulemaking regarding privacy, the FTC has instead proceeded by report, including, as noted above, the three reports produced in 2012. Those reports don't have to go through the demanding statutory processes and substantive filters that actual rulemaking would face, and that gives the FTC wide-ranging authority to create a zone of regulation beyond what it might create were the FTC to proceed through ordinary rulemaking. Instead, the FTC reports are an exercise in non-rulemaking making a type of non-rule rules which the FTC hopes will tilt the development of the online and mobile industries.

A. Regulating One-by-One through Ex Post Privacy Decisions

The FTC's adjudicatory privacy authority depends on an underlying framework provided by other law or by the regulated firms themselves. The FTC finds it easy to pursue companies for alleged deceptive privacy practices. The pattern in case after case is quite clear: a firm makes a statement about how it treats a consumer's information and then does something that the FTC sees as inconsistent with that statement. No matter how obscure the statement and seemingly without any reason to believe that any consumer has actually seen the statement or otherwise relied on it, the FTC is prepared to treat that as a deceptive practice and faced with that threat, the FTC enters into a consent agreement with the target.

Given that structure, firms might be wise not to make privacy statements at all, but underlying federal and state law may require them to do just that. A number of federal statutes, such as Gramm-Leach-Bliley and HIPAA require upfront privacy

[9] The so-called trade regulation rules are set forth at 16 CFR Chapter I, Subchapter D.

disclosures.[10] California has recently moved to expand the domain of the California Online Privacy Protection Act of 2003, which became effective on July 1, 2004.[11] California requires a conspicuous privacy policy from any operator of a commercial web site or online service that collects personally identifiable information about California individuals through the Internet.[12] That privacy policy must identify the categories of personally identifiable information collected and third-parties that that information might be shared with and must also set forth a mechanism for notifications of changes to the privacy policy. In February, 2012, the Attorney General of California announced a joint statement of principles to apply the California disclosure law to online app marketplaces.[13] On December 6, 2012, California sued Delta Airlines for failing to comply with its alleged obligation to provide a conspicuous privacy notice in its mobile app.[14]

So even though the FTC itself doesn't require general privacy statements by firms, other U.S. law may do exactly that. The consequence of that is that, given the FTC's understanding of its deceptive practice authority, simple mistakes by firms can be turned into broad regulatory interventions by the FTC for the

[10] See Federal Trade Commission, Bureau of Consumer Protection, Division of Financial Practices, The Gramm-Leach-Bliley Act Privacy of Consumer Financial Information (online at http://www.ftc.gov/privacy/glbact/glboutline.htm); U.S. Department of Health & Human Services, Health Information Privacy (online at http://www.hhs.gov/ocr/privacy/hipaa/understanding/summary/index.html).

[11] See http://www.privacy.ca.gov/privacy_laws/index.shtml#online.

[12] California Business and Professions Code Sec. 22575(a) (online at http://www.leginfo.ca.gov/cgi-bin/displaycode?section=bpc&group=22001-23000&file=22575-22579).

[13] Joint Statement of Principles by California and Amazon, Apple, Google, Hewlett-Packard, Microsoft, and Research In Motion, February 22, 2012 (online at http://oag.ca.gov/system/files/attachments/press_releases/n2630_signed_agreement.pdf).

[14] See State of California Department of Justice, Office of the Attorney General, Attorney General Kamala D. Harris Files Suit Against Delta Airlines for Failure to Comply with California Privacy Law, December 6, 2012 (online at http://oag.ca.gov/news/press-releases/attorney-general-kamala-d-harris-files-suit-against-delta-airlines-failure).

target firms. The FTC's 20-year settlements with Google, Facebook and mySpace are prominent examples of that approach at work, but less visible examples are also telling.

How does installing p2p software to download music turn into a deceptive practice prosecuted by the FTC? In two recent cases, the FTC alleged that two firms mismanaged their computer networks by allowing peer-to-peer software to be installed on those networks.[15] The one case involved a Toyota dealership in Georgia, the second a debt collector in Utah. In both cases, someone at the firm installed p2p software presumably to do any of the standard things that people do with that software. In the debt collector case, the software was installed by the firm's chief operating officer, and, as the FTC's complaint makes clear, there was no business reason for installing the software.[16]

There is no suggestion in the FTC complaints that the p2p software was installed with the intent of somehow making consumer information available throughout the p2p network. The most natural guess is that there was simply a complete disconnect between the process of installing the p2p software and the issue of how consumer information should be protected. A typical p2p user may have very little sense of the full implications of using the program and would have no reason to think that the firm's information was somehow at risk for being exposed to a much wider audience.

Focus on what is missing in this analysis. There is obviously no attempt to establish some sort of causal connection between the consumer behavior, say buying a car, the privacy disclosure and the alleged deceptive practice. Did the consumer rely on the disclosure in any way in entering into the transaction? Did the consumer even see the disclosure? If the consumer was presented with a required

---

[15] Federal Trade Commission, FTC Charges Businesses Exposed Sensitive Information on Peer-to-Peer File-Sharing Networks, Putting Thousands of Consumers at Risk, June 7, 2012 (online at http://www.ftc.gov/opa/2012/06/epn-franklin.shtm)./

[16] Federal Trade Commission Complaint, In the Matter of EPN, Inc. (online at http://www.ftc.gov/os/caselist/1123143/120607epncmpt.pdf).

form for signature, did the consumer read the form? Switch to the firm in question. There is no effort to establish a real sense of mens rea on behalf of the firm and so simple mistakes in execution are turned into vast schemes of deception. There is no operating space between mistake and deception.[17] The FTC doesn't have general authority to regulate the information security efforts of firms, but the combination of broad ex ante privacy policies coupled with the FTC's current use of its deceptive practice authority has effectively given the FTC the ability to implement rules one-by-one after the fact.

Were that structure not troubling enough, matters get even worse when we consider the broad authority that the FTC has to mash together different theories of liability and remedies. There is a very odd dynamic at work in consent orders between the FTC and the firms it targets. The settling firm will first and foremost internalize the possible remedy that it will face going forward as embodied in the consent order. That is the new individualized, company-specific legal regime that it will operate under. The settling firm cares much, much less about the liability theories set out in the complaint. It too will have to live with whatever value the complaint and its settlement have as precedent but most of the consequences of whatever theories of liability are floated in the complaint will be borne by others.

And if the firm can buy a narrower consent order at the price of a broader complaint is should almost certainly do so. That is something very much in the control of the FTC. The critical issue here is the fundamental non-linearity of causes of action and agreed-to remedies. If the remedy is going to the same regardless of whether one, two or three different causes of action are alleged, the FTC has a broad almost unilateral ability to build up liability theories through the consent decree process with very little reason for settling firms to push back in that process. In the absence of court litigation—and there has been none so far, though the

[17] [what does the legislative history look like on deception?]

pending Wyndham litigation may change that[18]—the FTC has been able to construct a data security common law through its consent orders. But this common law is mainly a one-sided common law, as targets have little reason to contest liability theories absent direct changes in the remedies they would face. And there is little doubt that lawyers advising clients look to this common law to guide their actions.[19]

In the U.S., we have something of a kludged together privacy regime. We use privacy disclosures as one way of informing consumers about the privacy implications of their choices, but, especially online, there are good reasons to think that privacy policies and the like actually do little to inform many consumers. Consumers need not read the policies, many of which, of course are quite long and written by lawyers, and the general learning of online contracting is that consumers spend almost no time with these documents.[20]

We shouldn't think of privacy policies as creating disclosures to consumers. They are really effectively disclosures to the government and privacy policies matter not because of what they tell consumers or because of how they alter consumer behavior directly but rather because of the ex post regulatory opportunities they create in favor the government, especially the Federal Trade Commission. The FTC's ability to declare as a deceptive practice any gap between what the privacy policy says and how it is executed means that the FTC has many chances to regulate the privacy practices of individual firms.

The core problem with the calls for much greater transparency in data collection and use is precisely the way in which the FTC's

[18] See Federal Trade Commission, FTC Files Complaint Against Wyndham Hotels for Failure to Protect Consumers' Personal Information, June 6, 2012 (online at http://www.ftc.gov/opa/2012/06/wyndham.shtm).

[19] See David Alan Zetoony, The 10 Year Anniversary of the FTC's Data Security Program: Has the Commission Finally Gotten Too Big for its Breaches?, 2011 Stan. Tech. L. Rev. 12.

[20] Cite M-W papers and others.

authority translates simple mistakes into actionable deceptive practices. The mistakes confer on the FTC the ability to regulate firms one-by-one without having the burden of establishing broader democratic legitimacy for the terms that it implements. It also means that mistakes that would otherwise be irrelevant are translated into regulatory differences across competitive firms.

B. FTC Privacy Rulemaking?

Switch to rulemaking. Rulemaking would be precisely the way in which the FTC could seek to implement a broader-based privacy regime. If you are opposed to more privacy regulation, you almost certainly are against an FTC rulemaking here. But lawmaking through consent decree and adjudication has its own problems. The party opposing the FTC in a particular situation may not have the resources required to enter into a serious contest with the FTC and, as noted above, may not have the same interests as other participants in the industry. A rulemaking proceeding would bring all of those resources and interests to bear together.

Over time, there has been uncertainty about the precise scope of the FTC's rulemaking authority. In 1971, when the FTC issued rules that required gasoline stations to post minimum octane ratings, the National Petroleum Refiners Association challenged the authority of the FTC to issue those rules. At that time, the FTC had acted under Section 6(g) of the FTCA, which was the general rulemaking authority conferred on the FTC in the original 1914 legislation.[21] The FTC had promulgated the octane rule as, as it were, an advance interpretation of a practice that it believed gave rise to an unfair method of competition or as an unfair or deceptive act or practice. The lower court concluded that the FTC did not have the authority to issue such trade regulation rules, but

[21] "That the commission shall also have the power— … (g) From time to time to classify corporations and to make rules and regulations for the purpose of carrying out the provisions of this Act."

the D.C. Circuit reversed, concluding that the FTC's rule-making authority should be interpreted liberally.[22]

But with that baseline in place, Congress responded with a series of statutory changes, each of which has made the exercise of the FTC's rulemaking authority more complex. In 1975, Congress enacted the Magnuson-Moss Warranty—Federal Trade Commission Improvement Act, which added a new provision, 15 USC 57a, regulating rulemaking proceedings for unfair or deceptive acts or practices. More changes were added in 1980 when a new Section 57b-3 was added to push the FTC towards cost-benefit analysis.[23]

And in 1994, Congress added a new substantive limit on the FTC's core Section 5 unfair practice authority and on its corresponding Section 57a rulemaking authority. The added 15 USC 45(n) bars the FTC from acting on an unfair practice claim "unless the act or practice causes or is likely to cause substantial injury to consumers which is not reasonably avoidable by consumers themselves and not outweighed by countervailing benefits to consumers or to competition." The new subsection then adds a second limit on the FTC's unfair practice powers: "In determining whether an act or practice is unfair, the Commission may consider established public policies as evidence to be considered with all other evidence. Such public policy considerations may not serve as a primary basis for such determination."

Taken together, these changes have made it harder for the FTC to issues actual rules, both procedurally and substantively. And the amendments themselves have highlighted the different flavors of FTC authority, as they seem to require fine distinctions among the three different substantive predicates for FTC action, namely "unfair methods of competition, "unfair … acts or

[22] See National Petroleum Refiners Ass'n v. Federal Trade Commission, 482 F.2d 672 (D.C. Cir. 1973).

[23] Pub. L. 96-252, May 28, 1980.

practices" and "deceptive acts or practices." The Section 5(n) limit applies only to unfair acts or practices, while the rulemaking changes added in 1975 and 1980 apply to rules for both unfair or deceptive acts or practices. None of the changes since *NPRA* have purported to alter the FTC's authority over unfair methods of competition. Of course, the original 1914 FTCA only referred to unfair methods of competition and it was only in 1938 that the statute was expanded to include unfair or deceptive acts or practices.[24] That means that the FTC's power to issue rules and the path for doing so depends on whether the rules in question purport to address unfair methods of competition or, instead, unfair or deceptive acts or practices.[25]

C. Privacy Regulation via Report: Non-Rulemaking

Given all of that, it is hardly surprising that the FTC might prefer to spend its nonadjudicative resources otherwise and it did that in 2012 in issuing its three interestingly different privacy reports. The March 2012 Privacy Report is a document for privacy professionals: a main report of 73 pages and 365 footnotes followed by three appendices, including a dissenting statement by Commissioner Rosch. Commissioner Rosch's dissenting statement from the report very much focused on the way in which the report was inconsistent with the Section 5(n) limit added in 1994.[26]

The October 2012 Facing Facts FTC staff report on facial recognition technologies purported to set forth best practices for a just developing technology. Commissioner Rosch again dissented from the report and revisited the ideas he had articulated in his dissent from the March privacy report. The report itself appears to try to navigate a tight line in what it does: "The recommended best

[24] Cite 1938 statute.

[25] The FTC's Operating Manual draws this distinction as well. See Ch. 7: Rulemaking at 7.4 (p33) (online at http://www.ftc.gov/foia/adminstaffmanuals.shtm).

[26] See at p. C-5 ("That is not only paternalistic, but it goes well beyond what the Commission said in the early 1980s that it would do, and well beyond what Congress has permitted the Commission to do under Section 5(n).").

practices in this report are intended to provide guidance to commercial entities that are using or plan to use facial recognition technologies in their products and services. However, to the extent that the recommended best practices go beyond existing legal requirements, they are not intended to serve as a template for law enforcement actions or regulations under laws currently enforced by the FTC. If companies consider the issues of privacy by design, meaningful choice, and transparency at this early stage, it will help ensure that this industry develops in way that encourages companies to offer innovative new benefits to consumers and respect their privacy interests."[27]

Take all of that in. So long as companies embrace the three core non-rules promulgated by the FTC in the March, 2012 privacy report—privacy by design, choice and transparency—all will be fine. Just apply those ideas to facial recognition technology. And if we, the FTC, have said something that isn't actually law, we recognize that we can't enforce those as law. Another exercise in non-law law. This is a kind of regulatory sleight of hand—dare I say a regulatory deceptive practice?—in that reports aren't rules and aren't obviously legally binding in any meaningful sense and the FTC seems to recognize that, but then what do we have?[28] Non-rule rules and non-law laws?

But it is the September, 2012 Mobile App guidelines that are the most troubling in this regard. This is a short document, barely five pages of text with big print and lots of white space. It seems clear that the FTC hopes that actual app developers will sit down and read the document and make product design choices based on the document. I assume that most teenage developers don't start writing their mobile app by reading the FTC Mobile App report, but if they did, they might easily conclude that the FTC had

[27] Xxx.

[28] Indeed, 15 USC 57b-3 expressly carves out from the definition of the "rule[s]" to which the 1980 amendments apply "interpretative rules, rules involving Commission management or personnel, general statements of policy, or rules relating to Commission organization, procedure or practice."

imposed on them a legal obligation to develop their apps with built with the principle of privacy-by-design.[29]

Of course, the FTC hasn't done anything of the sort. The FTC starts by offering guidance on how app developers can "comply with truth-in-advertising standards and basic privacy principles." Lawyers look for footnotes but this is a footnote-free document. The FTC's core regulatory authority is written remarkably broadly—"unfair methods of competition in or affecting commerce, and unfair or deceptive acts or practices in or affecting commerce, are hereby declared unlawful"[30]—plus the FTC has separate authority over false advertisements.[31] That said, nothing in the FTC's Section 5 powers addresses privacy directly, or, to be clear, any other particular product attribute.

So the reference to truth-in-advertising is clear enough, but it isn't clear where the FTC would get the power to specify "basic privacy principles." The mobile app guideline then moves on. It makes the point that there is no infant industry-style protection for small app developers, as both the big and the small have to comply with the law: "Laws that apply to establish businesses apply to you, too, and violations can be costly." It is hard to disagree with that, but the tight textual linkage between truth-in-advertising standards and basic privacy principles and the follow-on suggestion that violating the law can be costly seems to be intended to suggest to an app developer that failing to comply with either the advertising rules or basic privacy principles can give rise to a law violation.[32]

[29] Consider even statements by privacy professionals on this. See Jules Polonetsky & Omer Tene, It's Not How Much Date You Have, But How You Use It: Assessing Privacy in the Context of Consumer Data Integration, [date] ("The FTC requires companies to 'promote consumer privacy throughout their organizations and at every stage of the development of their products and services.'")

[30] 15 U.S.C. 45.

[31] 15 USC 52.

[32] Could test this with a survey on Amazon Turk.

After discussing truthful advertising, the mobile app guide turns to privacy. App developers are instructed to "[b]uild privacy considerations in from the start," meaning, as the next sentence makes clear, "privacy by design." And what does that mean: "[i]ncoprorating privacy protections into your practices, limiting the information that you collect, securely storing what you hold on to, and safely disposing of what you no longer need." Again, the FTC hopes to bake privacy-by-design into industry development practices. It is attempting to do that without promulgating actual rules that would comply with the procedural and substantive statutory standards for rule-making. And in documents like the FTC's Mobile Apps report, it is doing that in a way that seems like that can't possibly be transparent to the FTC's target audience of developers.

Governments that embrace the new tools of communication such as blogs and Twitter as the FTC has—to its credit—are in an awkward posture. On the one hand, the FTC is being accessible and not burying law in obscure publications like the *Federal Register*. FTC blog posts and reports like the Mobile App Guidelines are intended to be read by a wider audience, but seeking a wider audience comes with its own set of burdens. I suspect that if the FTC brought to bear its own deceptive act or practice framework to something like the Mobile Apps report, it wouldn't make the cut. There is too much of a gap between the behavior that the FTC is seeking to encourage versus what we know that the FTC can actually require, but that gap isn't at all apparent in the materials facing developers.

II. A Competitive Data Collection Market

The FTC's privacy by design and security by design standards need to be evaluated under the standards of proof set forth in Section 5(n). That section requires that the act or practice "causes or is likely to cause substantial injury to consumers which is not reasonably avoidable by consumers themselves and not outweighed by countervailing benefits to consumers or to competition." We

need a framework in which to make those assessments and I think that requires us to understand what we think a competitive data collection market would look like. And it is hard to think about what that market would look like without understanding the legitimate boundaries of data observation and storage. We can't think of the consumer as exchanging data for something if the firm on the opposite side has the independent right to acquire that data without a consumer's permission. There is no need for exchange—for quid pro quo—if the firm can just acquire the information on its own.

A. Observation and Joint Data Creation

A concrete example might be useful and a recent controversy involving Orbitz, the online travel agent, will frame this nicely. A *Wall Street Journal* investigation concluded that Orbitz was returning different responses to users based on the type of computer that they used to visit Orbitz's site.[33] Based on patterns that Orbitz had derived from observing how consumers interacted with its website, Orbitz had concluded that Mac users spend more on hotels than PC users and that it therefore made sense to return nicer hotels to Mac users. As we might expect, Orbitz's interaction with its customers generates an enormous amount of data: 750 terabytes of info of these interactions in a single year (with, according to Orbitz, one terabyte equaling 285 million pages of text).[34]

Start with the practice itself. The job of a travel agent is to pull out of the large universe of potential answers those that somehow match the consumer best. If it turns out that Mac users are

[33] Dana Mattioli, On Orbitz, Mac Users Steered to Pricier Hotels, The Wall Street Journal, August 23, 2012 (online at http://online.wsj.com/article/SB10001424052702304458604577488822667325882.html).

[34] Barney Harford, Orbitz CEO, Orbitz: Mac users book fancier hotels than PC users, USA Today Travel, May 9, 2012 (online at http://travel.usatoday.com/hotels/post/2012/05/orbitz-hotel-booking-mac-pc-/690633/1).

systematically different from PC users—and Apple has been telling us that for years—then we could easily imagine that different answers should be returned to Mac and PC users. That might turn out to be a relatively efficient matching rule and indeed one that might be seen as having the virtue of not being tied to the particular identity of the consumer in question. Think of this as an intermediate second-best matching proxy: not as good, perhaps, as a situation where Orbitz has installed a cookie on the particular user's computer so that Orbitz can instantly reflect all of the prior transactions of that user in choosing which results to return to the searcher but superior to having Orbitz just return the same undifferentiated responses to all consumers. And for the cookie-averse, this might be a good matching/privacy middle ground.

Be clear on the mechanics of what happened here. Orbitz acts as a travel agent to sell tickets to consumers. When a consumer engages on a transaction with Orbitz, Orbitz has the chance to observe and record much information about that transaction. I am hard-pressed to understand how we would say that one party or the other to the transaction somehow has a greater ownership interest in the raw facts of the transaction. Subject to constitutional limits such as the First Amendment, we can choose to regulate this, but I don't have a strong pre-law sense of ownership rules for the data in these transactions. We constantly interact with people in ways that those joint interactions give rise to data accessible to those directly involved in the interactions and often, some times uncomfortably against our will, to bystanders as well. If you and I enter into a business deal, I don't see an obvious basis for somehow assigning a greater set of rights to the raw facts of that deal to one person in the transaction or another.

We obviously have some settings where we explicitly establish a different legal framework for these interactions, but those are exceptions, not the rule. So conversations between lawyers and clients, doctors and patients, priests and penitents and marital partners are often situated differently with regard to what rules control what emerges from those interactions. The fact that we

have created separate regimes for these transactions emphasizes the underlying symmetry that exists in most joint interactions. And, in other contexts, certain characteristics are simply taken off of the table as a legitimate basis for decision-making. U.S. law forbids discrimination for a variety of characteristics, such as race and gender, but so far at least, there are no laws barring decisions based on computer type.

The takeaway point here is that we can't really begin to think about a competitive data market—a market in which consumers exchange their data for value from producers—without an understanding of what data producers can access without the permission of the consumer. Information that producers can observe and collect without permission is information that they need not exchange value for, though of course consumers can walk away, physically or virtually, and choose not to interact with a seller at all.

B. Mechanisms of Data Competition

Return to the competitive data market framing and the question of what we should expect in exchange-based transactions. On the analysis above, consumers can't expect to control what Orbitz can observe from consumer visits to the Orbitz website, so consumers can't engage in a direct quid pro quo data transaction with Orbitz. But the legal framing of data acquisition and ownership may be quite different from the competitive framing. Consumers of course need not visit the Orbitz website at all and there are competing websites vying for a consumer's business and the chance to gather data from the consumer as well. It is that competition which is of interest to us.

That competition could take on any number of forms and that is likely to depend on how much heterogeneity there is among consumers in their desire for privacy or for not having decisions as to them based on personalized characteristics. Consider three approaches: voluntary data collection limits; personalization signals; and consumer attraction incentives.

1. VOLUNTARY DATA COLLECTION LIMITS

Sites can compete on privacy by not collecting information. For example, DuckDuckGo.com, a competitor with Google in the general search market, states simply that it "does not collect or share personal information."[35] DuckDuckGo takes a variety of steps to prevent what it labels "search leakage," which is the transmission of search terms that are visited from a search engine. DuckDuckGo doesn't log header information transmitted to it from a user's browser, meaning that DuckDuckGo could not use Orbitz's personalization strategy. DuckDuckGo also doesn't not collect search histories so as to make sure that it has nothing to deliver to law enforcement should they issue a request for information to DuckDuckGo.[36]

2. PERSONALIZATION SIGNALS

Part of the unhappiness with Orbitz's practices was based on the idea that consumers received no notice of the fact that they were receiving personalized results.[37] Firms could provide a personalization signal, perhaps in English or an icon, to indicate to consumers that they were receiving personalized results. Some firms, such as Google, do this voluntarily. Over time, Google has moved into personalized search and it does that in different ways depending on whether you are logged into Google through your Google account or whether it is merely identifying your device through a cookie.[38]

---

[35] DuckDuckGo Privacy Statement (online at http://duckduckgo.com/privacy.html).

[36] Id. ("DuckDuckGo takes the approach to not collect any personal information. The decisions of whether and how to comply with law enforcement requests, whether and how to anonymize data, and how to best protect your information from hackers are out of our hands. Your search history is safe with us because it cannot be tied to you in any way.")

[37] Ed Bott, What Orbitz' sneaky web tracking is your problem, too, ZDNet, June 26, 2012 (online at http://www.zdnet.com/blog/bott/why-orbitz-sneaky-web-tracking-is-your-problem-too/5159).

[38] See Google, Basics: Search history personalization (online at

In early 2012, Google expanded search personalization by linking search results to Google+, Google's social product, in an effort to integrate social signals into search results.[39] Google may think that this is a natural evolution in search and the fact that it may also boost the presence of Google+, which competes with Facebook and other social sites, is just a convenient fringe benefit. The other possibility is that the linking is more of an effort to leverage Google's strong position in search to boost a much weaker position in social. But the issue here is the personalization signal. Google does this in English at the top of the organic search listing ("40 personal results. 1,100,000 results.") but also by adding a head-and-shoulders icon to the left of particular results to indicate that result is personalized. Google also provides paired icons at the top of the screen—another head-and-shoulders icon matched with a globe—and clicking on those icon makes it easy to toggle back and forth between personalized results and general results.

3. CONSUMER ATTRACTION INCENTIVES

Both data collection limits and personalization signals focus directly on data collection and use practices and privacy. But competition is more general and more important than those examples suggest. In the world of Orbitz, a sale to a consumer is about more than just about the ability to earn a cut on whatever is sold to the consumer. The consumer also brings data with her and the chance to observe that data is valuable. Firms will compete to attract consumers precisely to have the chance to collect that data. Consumers may not be able to withhold data from Orbitz or may not want to go through the effort to do so, but they can easily choose another site offering them more value and in so doing take their data elsewhere.

http://support.google.com/accounts/bin/answer.py?hl=en&answer=54041&ctx=cb&src=cb&cbid=15sk8nk13duco&cbrank=3); see also Google, Basics: Google Web History (online at http://support.google.com/accounts/bin/answer.py?hl=en&answer=54068).

[39] See Google Official Blog, Search, plus Your World, January 10, 2012 (online at http://googleblog.blogspot.com/2012/01/search-plus-your-world.html).

In this framing, it is critical to note that it isn't important to the analysis that Orbitz's data practices be transparent to the consumers. Data practice transparency matters for direct competition over data practices or over personalization signals. We should expect that type of competition to matter most for privacy-savvy or privacy-sensitive consumers. For the rest, we might think that firms would simply grab all of the data possible and that none of the value associated with that would flow to consumers. If you live in the world of Sec. 5 of the FTCA, you might call that unfair.

C. Framing Data Collection Harm

The data point is right, but the conclusion is wrong or at least partially so. If consumers aren't making decisions based on the volume or value of the data flowing to the firms that they are dealing with, firms have no reason to limit the amount of data grabbed, absent some separate regime that prevents that. But competition among firms for that valuable data will drive value to consumers. That could take the form of reduced prices for the underlying products or better services, but we shouldn't think, in equilibrium, that firms can just extract the data value from consumers without paying for it. The mechanism at work here, for the privacy-insensitive consumers, isn't the consumers themselves but rather the behavior of competing firms to attract those consumers. Those firms each understand the value of the data and recognize that they can't acquire the data unless the consumers choose to do business with them.

The main concern here is that firms will take data beyond the point at which an attentive consumer would want to transfer the data. This is a situation where consumers value not transferring the data more than firms value getting the data. In that situation, firms wouldn't buy the data from attentive consumers, as those consumers would demand more for the data than it would be worth for the purchasing firm to pay. For inattentive consumers, firms will just grab that data. To be sure, as noted above, they will compete against other firms to obtain that data, but the limit of the value transferred to consumers in that case will be the firm's value

of the data, or some fraction of that depending on the strength of the competition.

But, in the constructed scenario, that amount will necessarily be less than the actual value to consumers of not transferring the data in the first place. This problem shouldn't arise for attentive, privacy-focused consumers, but it will be a problem for inattentive consumers. This is the structure of the competitive failure that needs to be addressed, either through direct regulation of data transfer or by constructing a mechanism which reduces the costs of attention for the inattentive consumers, so as to get those consumers to engage more directly in controlling the amount of data they transfer.

Note that this is a pretty standard economic account of social harm. We think that we are seeing a version of market failure when a good is transferred systematically from an individual who values it more to an individual who values it for less. That isn't the direction in which we expect to see goods flowing in a healthy, competitive market. Now part of what is at stake here is the division of consumers into those who make decisions based upon privacy/data flow considerations and those who do not. If we believed that that split just reflected actual underlying values for privacy and data transfer in a world of heterogeneous consumers, that would be fine. Inattention to data transfers or privacy may be perfectly rational if you simply aren't concerned about the transfers or even believe that the tailoring of the transfers permit actually improves the experience for the consumer. The greater the extent to which we think that the data inattentiveness isn't rationally constructed or where we see consumers making data bets that turn out to be systematic losers after the fact, the more we should favor some type of intervention to shift the balance between attentive and inattentive consumers.

III. Perturbing Consumer Privacy Choices

The data collection competition discussion was framed on the idea of privacy-attentive and privacy-inattentive consumers. Privacy-

attentive consumers already may be making choices that reflect their concern about privacy or the relative data transfer practices of different services. Given the range of voluntary uses of privacy choice tools such as data collection icons and privacy signals, we should start by assessing data on what choices we are seeing by consumers. But we also need to figure out what we think we would like to see in equilibrium. Consumers ignore all sort of products characteristics in making many choices and there is no obvious reason to single out privacy as the characteristic that consumers some how must direct attention towards.

A. How Often Should Consumers Click on Ad Icons and Personalization Signals?

We need to figure out what rate of clicking we should hope to see on ad icons and personalization signals. One framing of this flows from the substantial number of surveys in which consumers indicate a preference against behavioral advertising.[40] If consumers really hate OBA as the surveys suggest, we should expect them to take opportunities to click on icons presented to them that would help them alter that experience, perhaps by opting out completely. On that framing, a meaningful gap between the number of consumers who had, say, opted out of OBA and those who express an aversion to OBA in surveys would be an indicator of the failure of the icon system to work well. And that gap might arise if the icon system was implemented in a limited way or if the icons themselves were confusing.[41]

We are between something of a rock and a hard place. On the one hand, if we place a great deal of stock in consumer surveys, then consumers don't seem to be implementing their desires

[40] For collected results, see Leon et al, What Do Online Behavioral Advertising Disclosures Communicate to Users?, CMU-CyLab-12-008, April 13, 2012.

[41] On this, see id; see also Komanduri et al, AdChoices? Compliance with Online Behavioral Advertising Notice and Choice Requirements, CMU-CyLab-11-005, October 7, 2011.

successfully.[42] On the other hand, surveys are just talk and actual behavior might be seen as a better indicator of what individuals actually care about. It is easy to care about something in the abstract, but the concrete may be more telling.

We might do better to consider the analysis in a particular setting, so focus on Google. On Google, ads typically appear at the top, side and/or bottom of the organic results. The top ads are separated by color—a palish tan—and are designated with a label "Ads related to [search]." That designation may be followed with an icon, the letter i in a circle. A click on that icon results in a pop-over box which in turn indicates that "these ads are based on your current search terms." The pop-over continues with more instructions: "Visit Google's Ads Preferences Manger to learn more, block specific advertisers, or opt out of personalized ads." On Google, the circle-i icon response is contextual. If Google believes that your search indicates that you are shopping for a particular type of product, you may received a sponsored shopping result in lieu of what would otherwise be the top-position side ads. In that case, a click on the circle-i icon results in a disclosure from Google: "Based on your search query, we think that you are trying to find a product. Clicking in this box will show you results from providers who can fulfill your request. Google may be compensated by some of these providers."

Google is a business built on metrics tied to counting and given its voluntary embrace of personalization signals, Google should have a great deal of information on how frequently consumers interact with these signals. And Google is constantly evolving its search tools and engages in frequent A/B testing to evaluate the relative worth of different search approaches. Google presumably could be using—and indeed may be using—precisely the same process regarding its personalization icons, but we do at some point need to articulate what result we should see in equilibrium.

[42] [What is the info on consumer opt out of ad personalization at Google or through the AdChoices program?]

I don't think that is obvious. Consumers who never click on ads have little reason to invest time in helping Google tailor ads to them. Consumers who are privacy-attentive may be unwilling to allow Google to tailor ads to them. Of course, Google could still accumulate information on a consumer but there is less reason to do so if the consumer doesn't want to receive personalized ads or services. Consumers who do click on ads might want to make that that screen space is used optimally, as, after all, the ads are different type of answer to a consumer's inquiry. At one point, Google reported that consumers who had opted in to receiving personalized ads were xx% more likely to click on those ads than consumers who had not done so. And, of course, the better that Google is at identifying my preferences the less reason I may have to interact directly with the ad preferences manager.

B. Tools for Perturbing Consumer Privacy Choices

The assorted data icon and personalization signal programs are very much an ongoing undertaking. Some observers see these programs as almost designed to fail, programs which arise in the face of threatened regulation and whose main goal is to defeat that threat. Privacy protection theater.[43] Others might see the usual difficulties of explaining complex products to busy consumers. But efforts to achieve greater transparency continue, consistent with the call for just that in the FTC's March 2012 report.[44]

Faced with this situation, the FTC could choose to do what it has done, which is to regulate firms one-by-one through its

[43] See, e.g., Lorrie Faith Cranor, Necessary But Not Sufficient: Standardized Mechanisms for Privacy Notice and Choice, 10 J. Telecomm. & High Tech. L. 273, xxx (2012) ("In the years that followed, we have seen a continuous cycle of new industry initiatives to improve notice and choice mechanisms and empower individuals, followed by a loss of interest in these initiatives when pressure from regulators subsides.")

[44] See Natasha Singer, A Tumultuous Trip to Mobile App Transparency, The New York Times, December 8, 2012 (online at http://www.nytimes.com/2012/12/09/technology/effort-to-clarify-mobile-app-data-rights-hits-snags.html?_r=0).

adjudicatory powers and to issue reports with non-rule rules. But there are other tools available to the FTC. Consider three ideas:

- *Mandated Data Disclosures and Icons/Personalization Signals*. The government often requires disclosure of information and the FTC has adopted this strategy in the past with rules requiring disclosure of gasoline octane rating and home insulation ratings. It can be quite difficult to know exactly what information is being accessed by a particular app.[45] Even something more basic, such as the difference between the bandwidth consumed by free and paid versions of apps, can be hard to access.[46]
- *Government Privacy Apps*. Mandated data disclosures suffer from the core problem that it can be hard to assess whether consumers actually value the information for more than it costs to create it. The government could go into the app business directly by creating a privacy dashboard app to give consumers a window into the data activity on their PCs or devices. The government would have the advantage that it could distribute the app for free. Consumers could download the app voluntarily, though of course, consumers might fear that the government app was part of the government's surveillance apparatus.
- *Regulatory Safe Harbors*. The FTC could tie its prosecution decisions to market results on how often consumers interact with a firm's privacy icons and signals. The adjudicatory structure of privacy for the FTC—statement followed by mistake treated as deception—must create many opportunities for the

[45] WSJ What They Know Series.

[46] See Ericsson Mobility Report, November, 2012, p.18-19 (free, ad-supported version of game app used 220 kB in a typical session while paid app version of same game used 1.3 kB).

government to prosecute firms. The FTC could adopt a policy of forbearance in circumstances where firms could demonstrate that their customers had been active users of the firm's privacy settings. This could incentivize firms to build effective privacy choice systems for consumers.

IV. By Design and The Standard of Proof for Unfairness

We are moving to a regime where privacy and security are dealt with through architecture and design. As noted above, in its recent reports, the FTC has emphasized the idea of "privacy by design" and in its enforcement proceedings it has invoked the related notion of "security by design." As the FTC articulates the idea of "privacy by design" firms "should promote consumer privacy throughout their organizations and at every stage of the development of their products and services."[47] As to security, in documents that the FTC clearly intends to provide guidance for developers, the FTC tells mobile app developers to "start with security."[48] And in commentary on its settlement with HTC, the FTC emphasized that the key step for a firm was to "build Security By Design into every aspect of their business."

We need to consider how privacy by design and security by design matches up with the demanding standards of Section 5(n). For the FTC to have the authority to declare an act or practice to unfair, the FTC needs to show that the act or practice "causes or is likely to cause substantial injury to consumers which is not reasonably avoidable by consumers themselves and not outweighed by countervailing benefits to consumers or to competition." Even articulating what the relevant act or practice is for this analysis may be difficult, but even if we can do that, it may be very difficult for

[47] FTC Privacy Report at p. 22.

[48] Federal Trade Commission Bureau of Consumer Protection Business Center, Mobile App Developers: Start with Security, Feb 2013 (online at http://business.ftc.gov/documents/bus83-mobile-app-developers-start-security).

the FTC to meet the statutory standard of Section 5(n) and it seems quite unlikely that the by-design approach meets that standard.

A. Privacy By Design

1. PRIVACY AND PAYMENT MECHANISMS

We should start by examining carefully the idea of privacy by design and we need to return to the discussion from above about the structure of data collection markets. We need a baseline transactional privacy regime to measure those markets against. I struggle to frame that. Should our baseline transaction should be a person wearing a shopping bag over his head paying with cash in a town far from his hometown? That seems silly, but perhaps that transaction minimizes data transfer to the seller and maximizes privacy for the consumer. This is an anonymous transaction done using cash. In the physical world, as soon as you introduce a payment mechanism other than cash, you are creating a paper trail for the transaction. Credit card transactions require that an identity-based relationship be established upfront between the consumer and the credit-card company. And the actual in-person transaction will require disclosure of the credit card and perhaps some other forms of information to establish identity.

Does privacy-by-design implicate the choice between fee-based content and advertising-based content? Fee-based online transactions also typically require the creation of an identity-based relationship. As we think about the important platform competition among Amazon, Apple, Facebook and Google part of what separates Amazon and Apple from Facebook and Google is precisely that Amazon and Apple have a vast database of credit card and other identity information. We could insist that online providers offering fee-based services arrange for prepaid cards that could be bought in the physical world for cash.

Prepaid cards are often used for gift-giving, but they would also enable the privacy hard-core to operate via cash online even without a laptop with a slot that accepts dollar bills. Executing the

prepaid cards may require the establishment of a temporary or even persistent identity with the service provider, such as iTunes, but that need not be the consumer's actual identity. And the identity certainly need not be a consistent identity of the sort that is likely to be tied to a credit card used both online and in the physical world. We could imagine prepaid cards that work across multiple providers.[49] That would avoid forcing each online provider to create its own card and network of vendors for distributing those cards. It seems hard to imagine that privacy-by-design is intended to require fee-based online providers to establish a payment regime that comes closest to mimicking how cash operates in physical space.

Switch to advertising-supported online services. Free content with unpersonalized advertising with no tracking and data collection seems maximally privacy protecting. Media-based advertising—advertising in newspapers, magazines, TV, radio and cable—might target particular demographic groups, but it was rarely, if ever, designed for a particular individual. Billboards were untailored as well. Direct-mail advertising might have come closest to being genuinely targeted and individualized.

Of course, online advertising wasn't initially personalized either. Google built its business off of contextual advertising and moved relatively late into interested-based—personalized—advertising.[50] Contextual advertising can be quite refined. Indeed, Google's ability to define a highly-textured content niches through long-tail searches is part of why it has succeeded with advertisers. None of that requires personalization in the sense of offering individuals running exactly the same search different ads.

[49] This would be similar to a reloadable debit card, though those cards often require the establishment of an actual identity with the card vendor. See Dana Dratch, 6 things to know about reloadable prepaid cards, CreditCards.com, October 27, 2011 (online at http://www.creditcards.com/credit-card-news/6-things-know-reloadable-prepaid_cards-1271.php). See also Walmart MoneyCard Reloadable Prepaid Card application (online at http://www.walmartmoneycard.com/walmart/getacardnow).

[50] [Need date].

What standard does a privacy-by-design framework impose on a step towards some sort of personalized advertising? Is it enough for the firm to contend that the firm would make more money if it engaged in personalized advertising? And if that isn't sufficient, what standard is going to operate instead? The hard-core version of privacy by design would seem to suggest a hierarchical approach to payment mechanisms: cash and cash substitutes such as gift cards; contextual advertising; personalized, anonymous advertising and perhaps only then identity-based cash payments (checks or credit cards).

Return to the standards of Section 5(n). Is the relevant practice that we should be considering through an unfairness lens online behavioral advertising itself?[51] Given Section 5(n), that would seem to require a square focus on the "countervailing benefits to consumers or to competition" from OBA. It would seem like a tall order to assess the benefits to consumers of OBA. Is the question to be asked what content or services are deliverable to consumers with OBA that could not be delivered in a world with say only contextual advertising?

In economies based on competition, we typically don't try to answer those questions in a centralized way. Section 5(n) looks to consumers but it also looks directly to benefits to competition. The essence of competition is precisely that we let the market sort out whether an innovation like OBA turns out to be a practice that adds incremental value. We don't hardwire the economy to insist that all services or content delivered to consumers needs to be paid for in cash or cash equivalents by those consumers. Traditional advertising in media substituted, in whole or in part, for cash payments that would otherwise have been required by consumers. And all of that is without considering the extent to which the

[51] In his dissenting statement from the March, 2012, FTC Privacy Report, Commissioner Rosch characterized the report as keyed to the unfairness rationale ("First, the Report is rooted in its insistence that the 'unfair' prong, rather than the 'deceptive' prong, of the Commission's Section 5 consumer protection statute, should govern information gathering practices (including 'tracking').) Rosch Dissenting Statement at p.C-3.

practice in question is "reasonably avoidable by consumers themselves" and how that standard should operate in a world of consumers with different values for privacy and data transfer.

2. DEFAULT SETTINGS IN DO-NOT-TRACK

The behavior of industry participants is almost always indicative of what is important and what is not. The fight over default setting in two Internet browsers, Microsoft's Internet Explorer 10 and an update to Mozilla Firefox is a clear statement about the importance of default settings for privacy. And the fight is a nice example of exactly how complex privacy by design can be in practice. On May 31, 2012, Microsoft announced that for Internet Explorer 10 on Windows 8, Microsoft would set "Do Not Track" as the default setting for the browser on installation.[52] Such an approach seems to track quite closely the FTC's privacy-by-design call. Microsoft clearly saw this as a competitive move designed to attract consumers to IE 10:

> This decision reflects our commitment to providing Windows customers an experience that is "private by default" in an era when so much user data is collected online. IE10 is the first browser to send a "Do Not Track" (DNT) signal by default. … While some people will say that this change is too much and others that it is not enough, we think it is progress and that consumers will favor products designed with their privacy in mind over products that are designed primarily to gather their data.[53]

[52] The draft specification offers three different default settings for the do-no-track signal: off, on and no consumer choice indicated. See Lorrie Faith Cranor, If you choose not to decide, your web browser will make your choice, TAP Blog, June 3, 2012 (online at http://www.techpolicy.com/Blog/June-2012/If-you-choose-not-to-decide,-your-web-browser-will.aspx).

[53] IEBlog, Windows Release Preview: The Sixth IE10 Platform Preview, May 31, 2012 (online at http://blogs.msdn.com/b/ie/archive/2012/05/31/windows-release-preview-the-sixth-ie10-platform-preview.aspx).

It is hard to know what to make of that. Perhaps Microsoft bought into the FTC religion of privacy by design, and indeed, subsequent Microsoft statements invoked that framework.[54] Perhaps, more cynically, Microsoft understood the landscape of the competitive space in privacy. Information is important to many Microsoft competitors, but much less so to Microsoft. Microsoft has faced antitrust challenges in both the U.S. and the EU for tying Internet Explorer to Windows, but it continues to have the right to do that in the U.S. and with its do-not-track default for IE10, Microsoft was using that leverage to change the competitive playing field in privacy and data collection.

The Microsoft announcement led to a flurry of responses from competitors, advertisers and other software developers. Mozilla Firefox announced that it would take a different approach. Browsers could be set with three different default settings: accepting tracking, blocking tracking or with no indication of a user choice. Microsoft had announced that it would block tracking, but Mozilla was going to leave the choice entirely to users and not set an acceptance or a rejection default.[55] This is exactly the sort of competition in privacy that society should want.

The advertising industry immediately asked Microsoft to reverse course. The Association of National Advertisers (ANA) announced its strong opposition to Microsoft's privacy-by-design implementation.[56] On behalf of its 450 members, the ANA asked

[54] Brendon Lynch, Microsoft Chief Privacy Officer, Do Not Track in the Windows 8 Setup Experience, Microsoft on the Issues, Aug 7, 2012 ("Our approach to DNT in IE10 is part of our commitment to privacy by design and putting people first.") (online at http://blogs.technet.com/b/microsoft_on_the_issues/archive/2012/08/07/do-not-track-in-the-windows-8-set-up-experience.aspx); Brian Prince, Microsoft IE 10 Do Not Track Default Setting Assailed by Advertisers, eWeek, Oct 3, 2012 (online at http://www.eweek.com/cloud/microsoft-ie-10-do-not-track-default-setting-assailed-by-advertisers/).

[55] See Alex Fowler, Do Not Track: It's the user's voice that matters, Mozilla Privacy Blog, May 31. 2012 (online at http://blog.mozilla.org/privacy/2012/05/31/do-not-track-its-the-users-voice-that-matters/).

[56] ANA Strongly Opposes Microsoft's Decision Regarding Internet Explorer 10 'Do Not Track' Function, June 1, 2012 (online at

Microsoft to reset IE10's DNT signal to a default of off, so that consumers would have the chance to experience interest-based advertising and not suffer the horror—my phrase, not theirs—of "untargeted, irrelevant online advertising." The ANA Board of Directors followed up in October, 2012 with a letter to Microsoft to express its "profound disappointment with and strong opposition" to Microsoft's approach to DNT.[57] The Digital Advertising Alliance made clear that it would not require its members to respect what it termed "browser-manufacturer choice" rather than actual consumer choice.[58] Of course, there is no good way for a website to tell if the DNT flag was set to "on" by default by the browser maker or by direct choice by the consumer.

The obvious question is how we should think about implementing a default setting. Privacy-by-design seems to mean just that: establish a default setting that maximizes privacy for the consumer and force the consumer to affirmatively change that setting. An alternative would be to rely on survey data and to implement the views seen there. Of course, how surveys are framed matters, but one prominent behavioral advertising survey concludes that consumers do not want to received tailored ads.[59] The ANA board instead focused on what it saw as the collective action problem for advertising supported media: consumers want the content and want someone else to be exposed to the ads. The ANA put IE's market share at 43% and suggested that that large a drop in consumer's participating in behavioral ads would mean a real loss of free Internet content.

http://www.ana.net/content/show/id/23613).

[57] Association of National Advertisers, ANA Board Opposes Microsoft's Decision to Implement 'Do-Not-Track' Default Function for Internet Explorer 10 Browser, October 1, 2012 (online at http://www.ana.net/content/show/id/analetter-microsoft).

[58] Digital Advertising Alliance Statement, Digital Advertising Alliance Gives Guidance to Marketers for Microsoft IE10 'Do Not Track' Default Setting, October 9, 2012 (online at http://www.aboutads.info/blog/digital-advertising-alliance-gives-guidance-marketers-microsoft-ie10-%E2%80%98do-not-track%E2%80%99-default-set).

[59] [Turow 2009 study]; Hoofnagle?

As applied to Internet browsers, a requirement of privacy-by-design seems to call for a relatively clear result, and one that matches how Microsoft is approach IE10. The response of the advertising industry was perhaps to be expected. Their behavior suggests that they believe that many consumers are unlikely to override whatever default is set in the browser and therefore they see the IE10 default as taking away many opportunities to serve targeted ads. But Microsoft also faced condemnation from other parts of the Internet community. Roy Fielding, a co-author of the HTTP specification and a co-founder of the Apache server project, found Microsoft's responsive to be an abuse of the open standards process and rolled out a patch to Apache that would ignore the default setting set by Microsoft.[60]

Consider how privacy-by-design for browser default settings matches up with the standards of proof established in Section 5(n). On substantial consumer injury, presumably we would need to consider the possible privacy harms from the extra data that could be collected from tracking. We would need to weigh that against the benefits that consumers receive from the incremental free services and content supported by that tracking. And on avoidability and the consequences for competition, those seem to move together here. We seem to have relatively robust competition in Internet browsers and not mandating a particular setting for defaults—and privacy-by-design would seem to do just that—means that privacy becomes a more important aspect of that browser competition. That plays into avoidability: with meaningful choices among browsers, consumers can fully avoid the harm, if any, associated with a particular browser setting by choosing a different browser with a more-desired approach to privacy.

I assume that the real issue here is what to do about the privacy inattentive, as they are the consumers least likely to choose browsers based on privacy characteristics. One possibility is to do

[60] See Gavin Clarke, Apache man disables Internet Explorer 10 privacy setting, The Register, September 10, 2012 (online at http://www.theregister.co.uk/2012/09/10/fielding_apache_ie10_windows_8/).

nothing at all about them: inattention can be as much a choice, and sometimes a rational one, as lots of attention. But a different possibility available in networked spaces such as this one is that we could perturb the default browser settings for consumers or at least offer a pop-up screen presenting such as default setting opportunity. The EU did something like that as a remedy in its IE antitrust action against Microsoft, where EU consumers were to be presented with a browser choice screen the first time they used a Windows 7 PC.[61] Part of the choice on that has to be how much you think privacy inattention matters versus the costs of government engineering of computers.

B. Security By Design: The FTC's Cases against HTC

The FTC's recent action against smartphone producer HTC provides a nice context to talk through the by-design approach. On February 22, 2013, the Federal Trade Commission announced a new proposed consent order in an agreement that the FTC had reached with HTC. HTC is a smartphone leader, trailing only Apple and Samsung.[62] As described by the FTC in its press releases, blog posts and settlement documents, HTC "engaged in a number of practices that, taken together, failed to employ reasonable and appropriate security in the design and customization of the software on its mobile devices." Call this for short unreasonable security design. The story as set forth in the complaint gets technical quickly but we should make be clear to see the big picture.

1. COMPETITION IN SMARTPHONE PLATFORMS

The smartphone market is a platform market dominated by two platforms: Android, associated with Google, with, in March 2013,

[61] See www.browserchoice.eu.

[62] comScore, comScore Reports March 2013 U.S. Smartphone Subscriber Market Share, May 3, 2013 (online at http://www.comscore.com/Insights/Press_Releases/2013/5/comScore_Reports_March_2013_U.S._Smartphone_Subscriber_Market_Share).

a 52% market share and Apple's iOS, with a 39% share.[63] But Android and iOS are quite different in that Apple offers iOS as part of its vertically-integrated platform, while Android is an open source project.[64] That means that firms can create their own versions of Android and can use that customization as a basis for competition.

The underlying operating system for modern devices like the smartphone is quite different from the traditional structure of OSs for devices like personal computers. Networked devices raise different security issues than offline devices and the PC of course was originally designed as a pre-Internet device. Smartphones and tablets like the iPad are born networked—indeed, they are inconceivable as such without attached networks—and that means that security is an important part of the basic design of these devices.

That shows up directly in two related constitutive features of modern device operating systems, namely, sandboxing and permission models.[65] Sandboxing is a way to silo applications from each other and from the underlying functionality of the device. By default, applications don't interact with each other. If I install a new game on my smartphone, we might not think that it would have any need to access my contact information on the phone and sandboxing bakes that into the OS platform. But of course there can be many situations in which a smartphone application needs legitimate access to other information on the device or to the hardware functionality of the device and that is where permission model comes in. The permission model of a particular OS is exactly the way in which that platform organizes rights to use device functionality or device information.

---

[63] Id.

[64] See Android open source project (online at http://source.android.com/).

[65] The discussion here draws on K.W.Y. Au, Y.F. Zhou, Z. Huang, P. Gill and D. Lie. A Look at SmartPhone Permission Models. In *xxx*, date.

The computer science literature on permission structures for devices like smartphones make clear that there are real trade-offs that are made and that can be made in designing permission structures. These are interface choices in a context where "there is no consensus on the best way to design a permission system."[66] We should expect platforms to compete in their approaches to sandboxing and permission structures and indeed they do just that. The two leading smartphone OSs, Android and iOS, take very different approaches to their permission models and those are still different from the permission models of the Windows Phone OS or the Blackberry OS.[67] The Android permission model has changed rapidly between the first Android 1.0 release in September 2008 and the 2.3 release in December 2010[68] so any snapshot is somewhat arbitrary, but at one point, Android was offering 75 different permission settings to third-party applications at a point at which iOS was offering one such setting (user location). The Windows Phone OS was offering 15 such choices, Blackberry OS, 24, and Maemo, Nokia's successor to its Symbian OS, zero.

The point of that is the competing platforms take very different approaches to their permission models and there is no obvious reason we shouldn't think of that as an important point of competition. And, like all platform competition, the permission model is part of the multi-sided competition occurring on smartphone platforms. The permission model needs to mediate between the desires of application developers and those of smartphone users.

To see that, take just one aspect of permission model design, namely, when should the relevant permission request be made? Android devices typically present their permission requests at the

[66] See Adrienne Porter Felt et al, How to Ask for Permission, USENIX Workshop on Hot Topics in Security (HotSec) 2012 (online at http://www.eecs.berkeley.edu/~afelt/howtoaskforpermission.pdf).

[67] Table 1 of Au et al.

[68] Table 2 of Au et al.

point of the installation of the app. Does a long list of permissions unduly discourage consumers from installing an app? Is that a good thing, as it reduces the chances that an app developer will, as the OS community puts it, overdeclare, meaning seek more privileges than the application actually needs? Or do long permission lists at installation become EULA like in their numbing effect such that consumers ignore the list and just click forward? That would seem to make it much easier for developers to overdeclare.

Again, the point is here that the permission model is point of competition between different OS platforms. And that competition is even richer than I have suggested so far. The Android platform is open source, so by design it is possible for a firm to create its own version of Android and to compete with other Android devices by offering its own specialized version of Android.

2. PRE-INSTALLATION OF APPS ON SMARTPHONES

So far I have emphasized the legitimate way in which firms can compete in smartphone OSs, competition which includes different approaches to the permission models for using different tools of the smartphone. Again, this is multi-sided competition with developers, advertisers and consumers all having a stake in how the platform is organized. In looking at financing models for handsets like smartphones, we need to look at the role of preinstalled applications. Devices like tablets and smartphones are key ways for distributing software. Being able to preinstall software on the device is valuable and firms that control that—either device producers like Apple, Samsung and HTC or cell phone carriers—want to leverage that position.

The industry has a variety of names for these pre-installed applications—three, with all with varying degrees of attitude, are crapware, bloatware and shovelware[69]—but sophisticated observers

[69] Mike Jennings, Smartphone crapware: worse than laptops?, PC Pro, August 22, 2011 (online at http://www.pcpro.co.uk/blogs/2011/08/22/smartphone-crapware-worse-than-laptops/).

at least acknowledge the forces at work: firms with control over devices sell access to those devices and that wouldn't be worth as much if consumers could easily remove the pre-installed apps. The nature of two-sided competition is that a tax on one side of the platform may be competitively sensible if it buys more value on the other side of the platform. So the fact that consumers may find it irritating that they can't uninstall the Newstand icon from the iOS desktop doesn't mean that Apple is behaving stupidly—that seems unlikely of course—or even anti-competitively.

Preinstalled applications will of course be installed outside of the installation regime for consumer downloaded applications. They come with the device and are installed by the manufacturer or by the phone carrier. Again the choice of whether to have preinstalled apps or which apps to preinstall is a natural point of competition across handsets and across OS platforms. And consumers themselves may benefit on net, even if that isn't particularly salient to them. Consumers focus on the unremovable application and the fact that the overall price of the platform may have been lowered to make that possible is lost to them in the sands of time. The question is whether it worth more to someone on the other side to have a guarantee that the Newstand isn't going away than the cost to consumers of having an app that they can't remove.

3. The Introduction of HTC Sense

The FTC complaint against HTC is sufficiently opaque that some guesswork is required to provide a better context for the allegations. We should start with the state of the smartphone market at the end of 2008. Nokia was the worldwide leader in smartphone sales selling roughly 60.9 million units, followed by Research in Motion with 23.1 million units, Apple with 11.4 million units, HTC with 5.9 million units, Sharp with 5.2 million units and an undifferentiated group of others (including Samsung)

with 32.7 million units.[70] And those sales translated into roughly corresponding market shares for smartphone OSs: Symbian (the Nokia OS) was the market leader with a 52.4% share, followed by RIM with 16.6%; Windows Mobile with 11.8%; Mac OS X with 8.2%; Linux with 8.1%; and Palm OS with 1.8%. That list didn't include Android, as Android didn't even have an official release of its software developers kit until September 23, 2008.[71]

On June 24, 2009, HTC announced a new Android-based smartphone, the HTC Hero. The phone itself was fine—"a solid entry in the Android phone camp" according to one review—but what really was interesting was that HTC had created a new interface for the phone, called HTC Sense, which sat on top of the Android operating system. This was the first customized interface for an Android phone and made clear that the vision of competition within the Android ecosystem could be realized. And HTC Sense was seen as a play by HTC to create its own interface to its smartphones, as HTC was expected to bring HTC Sense interface to other phones, even if they were running a different OS, such as Windows Mobile.[72] Sense was certainly a hit with some of the tech crowd, who worried about whether the "sexy, wonderful UI that have people going gaga" was going to make to other HTC devices and who feared that Google and T-Mobile would block their users from Sense.[73]

[70] See Gartner Press Release, Gartner Says Worldwide Smartphone Sales Reached Its Lowest Growth Rate With 3.7 Per Cent Increase in Fourth Quarter of 2008, March 11, 2009 (online at http://www.gartner.com/newsroom/id/910112).

[71] Dan Morrill, Announcing the Android 1.0 SDK, release 1, Android Developers Blog, Sept 23, 2008 (online at http://android-developers.blogspot.com/2008/09/announcing-android-10-sdk-release-1.html).

[72] See James Kendrick, HTC Hero with Sense—First Android Phone with Customized Interface, gigaom.com, June 24, 2009 (online at http://gigaom.com/2009/06/24/htc-hero-with-sense-first-android-phone-with-customized-interface/).

[73] See Casey Chan, HTC Sense Coming to HTC Magic, Not T-Mobile myTouch 3G, AndroidCentral, Aug 17, 2009 ("We're not exactly sure how the legalities work but it probably has to do with Google and T-Mobile not wanting non-Google/T-Mobile approved software floating around on their handsets.") (online at

I am not sure how we would have assessed the state of competition in smartphones as of mid-2009 when HTC released Sense. We have the benefit of hindsight and know that there were dramatic shifts in market shares since then, both for handset makers with the ascendancy of Samsung and Apple and the collapse of Nokia and the corresponding rise of Android and iOS and the decline of Symbian and RIM. All of that suggests that this was a time for many competitive initiatives and that consumers have benefited enormously from that process. HTC Sense, which seemed like a possible winner in mid-2009, has to be understood as part of that competitive process.

4. HTC'S SECURITY DEFECTS

But HTC Sense wasn't perfect and the combination of Android with HTC Sense customization and pre-installed applications from cell carriers resulted in phones with security flaws. A blog post in June, 2010 discussing the HTC EVO argued that "[t]he Sprint customizations of Android are so bad that an Android application could get access to all of your data with very little work."[74] In October, 2010, the *Financial Times* reported that an investigation of the Android Kernel on the HTC Droid by security firm Coverity also demonstrated security flaws that would allow access to sensitive information on the device.[75] And by October, 2011, an Android news site reported what it termed a "massive

http://www.androidcentral.com/htc-sense-coming-htc-magic-not-t-mobile-mytouch-3g).

[74] See HTC EV) 4G: Nice Hardware, Horrible Sprint Software, grack.com, June 3, 2010 (online at http://grack.com/blog/2010/06/03/htc-evo-4g-nice-hardware-horrible-sprint-software/); Sprint.com, Find and update the software version of your HTC Hero (online at http://support.sprint.com/support/article/Find_and_update_the_software_version_on_your_HTC_Hero/case-gb746811-20091009-155624).

[75] Joseph Menn, Android faces critical security study, Financial Times, Oct 31, 2010 (online at http://www.ft.com/cms/s/2/10b955ba-e519-11df-8e0d-00144feabdc0.html#axzz2Sjfs5Hw3); Joseph Menn, Questions Arise About Google Android Security Risks, CNBC.com, Nov 12, 2010 (online at http://www.cnbc.com/id/40150849).

security vulnerability" in HTC phones running Android and HTC Sense.[76]

It is hard to know what to make of this. Software has bugs and has had them since the first line of code was written. We need more context to evaluate the bugs on the HTC smartphones. To take one example, the Coverity report referenced by the *Financial Times* was eventually made public. That report found what it termed 88 "high risk" defects in the version of Android tested on the HTC Droid Incredible and yet also concluded that the tested software "has approximately half the defects that would be expected for average software of the same size."[77] That sounds like real problems and yet better than average.

To take another example for context, consider in more detail the FTC complaint against HTC. HTC was alleged to created so-called permission re-delegation problems in pre-installing applications. As noted above, in the sandboxed world of smartphone apps, an app has to be given permission to access various system resources. Permission re-delegation occurs when an app without permission to access a particular resource is able to piggyback access through a second app that has permission to access the resource in question.

Permission re-delegation appears to be a widespread problem for Android applications, even in the underlying general Android distribution. One study found 15 permission re-delegation problems in five of the sixteen pre-installed applications required to be installed on any smartphone to meet the Android 2.2 compatibility definition.[78] As noted by one of the co-authors of that study in her slides on the study, this means that roughly one-

[76] Artem Russakovskii, Massive Security Vulnerability In HTC Android Devices, Oct 1, 2011 (online at http://www.androidpolice.com/2011/10/01/massive-security-vulnerability-in-htc-android-devices-evo-3d-4g-thunderbolt-others-exposes-phone-numbers-gps-sms-emails-addresses-much-more/);

[77] Coverity, Coverity Scan: 2010 Open Source Integrity Report, p.2.

[78] See A.P. Felt, H.J. Wang, A. Moshchuk, S. Hanna and E. Chin. Permission Re-Delegation: Attacks and Defenses. In xxx. Date.

third of the required pre-installed apps had permission re-delegation issues and these were apps written by Google's own Android engineers, quintessential insiders with deep knowledge of Android's internal operations. And stepping beyond Android's own pre-installed apps, the study found that more than one-third of 872 tested Android applications had re-delegation risks.

The nature of prosecutorial decisions by the FTC (and really any prosecutor) is that is can be almost impossible to assess why one case was brought and another possible case was ignored. As just suggested, security problems are widespread in software, even software written by individuals at the top of their profession, and nothing in the FTC action against HTC gives us any way of assessing whether the scope of problems alleged in the FTC complaint were distinctive as compared to other software. The FTC's own ever-so-brief public analysis of the complaint and consent order—slightly more than three pages single-spaced—certainly doesn't do any analysis of this sort.

5. How Much Should Firms Investigate Security Before Releasing Products?

We might summarize *HTC* this way: HTC set out to develop an innovative interface for its version of Android. That presumably was a substantial undertaking and exactly the sort of competition that benefits consumers. It made some mistakes in creating this interface, those mistakes were discovered, and it set out to correct them. HTC almost certainly could have invested more resources in trying to eliminate security mistakes upfront and a company organized around the idea of security by design might well have done that. But we don't know how the number of mistakes made by HTC compares to some sort of industry baseline and, perhaps more importantly, we don't have any sense of how a security-by-design requirement would have discouraged the overall level of innovation that HTC Sense represented. It is only in a framework in which regulatory burdens like security by design are seen as free lunches—as having only benefits and no costs associated with them—that we could eagerly embrace the regulatory remedy that

emerges from the HTC consent decree. As to that, the lesson that the FTC wants hardware and software designers to learn from the HTC case seems clear: "Whatever the cliché du jour, the message remains the same: Savvy companies build Security By Design into every aspect of their business."[79]

In truth, the right reaction is just the opposite: compared to the offline world of one-shot design choices, firms engaged in online processes should do less ex ante testing of products over dimensions like security and privacy. With networked, software-based products, firms can evolve those products in place and can fix problems in software as they are discovered naturally through use. In the offline world, once a product has been distributed, it has the characteristics it had when it left the factory floor and it can't readily be fixed after the fact. In contrast, with networked products, use by customers will reveal information about the product and that can be a good substitute for ex ante testing of the product.[80] With one shot-design, we can't readily substitute between ex ante testing and ex post testing—actual use of the product—but with networked products that can evolve in place, we can do just that.

[79] Lesley Fair, Batten down the patches: Six points to take from the FTC settlement with HTC, FTC Bureau of Consumer Protection Business Center Blog, Feb 27 2013 (online at http://business.ftc.gov/blog/2013/02/batten-down-patches-six-points-take-ftc-settlement-htc).

[80] Consider a stylized situation to see this. Suppose that a firm faces a design decision for a product. There is some uncertainty about that choice and the firm can incur a testing cost of C to determine the implications of the design choice. If the test is run, the potential problem can be eliminated, but absent the test, the firm will release a product that will result in a harm H with some probability p. In that circumstance, we would like the test run if the cost of the test is less than the expected harm ($C < pH$) and not otherwise. Now switch the situation and divided the harm H into $H = H1 + H2$. Again the harm will occur with probability p, but now, if the harm occurs, the firm can update the product in place and avoid the second chunk of the harm H2. Now we want the test run if the costs of the test is less than the unavoidable expected harm, or if $C < pH1$ and not otherwise. In situations in which $pH1 < C < pH$, we want the firm that can update the product to release the product without running the test even though a firm that could not update the product should run the test. We want less testing of products than can be updated in place, not more.

We should be clear on what was just said. For a given level of possible harm associated with a product and a given probability that that harm will arise, a firm should do less design testing in a networked world compared to an offline world because the firm will get free information about the status of the product through subsequent use and through the network has the ability to fix products in place. Harms that would be inevitable offline, if a bad outcome has taken place, can be eliminated after-the-fact for networked products through updates. Of course, there is a separate question of how the very fact that we are dealing with networked products changes the possible scale of harm that might result or the probability of harm, so the framing here shouldn't be thought of as networked vs. unnetworked products but rather about the role that updates play in altering the balance between ex ante learning about products through testing and ex post learning through use. And, to be clear, the FTC's analysis in the HTC cases doesn't seem to make anything of how networking might change possible harms.

Conclusion

The Federal Trade Commission has latched on to its broad Section 5 authority over unfair or deceptive acts or practices to jump in to regulate privacy and data security with an entrepreneurial nimbleness worthy of the private sector. That is what happens when governmental vision and rapid technological change intersect with an open-ended statute. But the FTC has largely been able to act in an unchecked fashion. The parties that it targets find it advantageous to negotiate consent decrees and those targets care most about the remedies that they face under those decrees and much less about the theories of liabilities alleged. That dynamic has made it possible to build up a common law of privacy and data security and to move beyond a deceptive practice framing towards one tied more directly to unfairness.

In doing that, the FTC has acted outside the statutory confines of Section 5(n) of the statute, which limits FTC authority over

unfair or deceptive acts or practices to a well-defined set of facts requiring substantial reasonably avoidable consumer injury without countervailing benefits to consumers or competition. The FTC's current doctrine, as defined not in rulemakings but instead through reports and not through actual adjudication but instead through consent decrees, doesn't confront those limits in a serious way and that is seen most directly in its recent emphasis on privacy by design and security by design. These amount to efforts to create broad restrictions on the processes that firms use to develop products and to do so without any sense of the trade-offs that firms make in designing products or of the trade-offs that consumers value in products.