New York University
A private university in the public service

School of Law
Faculty of Law

40 Washington Square South, Room 326
New York, NY 10012-1099
Telephone: (212) 992-8909
Facsimile: (212) 995-4760
E-mail: ira.rubinstein@nyu.edu

Ira S. Rubinstein
Adjunct Professor of Law
Senior Fellow, Information Law Institute

Comments of Professor Ira Rubinstein on
Privacy by Design and the FTC's Proposed Framework

Re:

Federal Trade Commission (Bureau of Consumer Protection),
Preliminary FTC Staff Report on "Protecting Consumer Privacy in an Era of Rapid Change: A Proposed Framework for Businesses and Policymakers" (December, 2010)

February 18, 2011

Introduction

My name is Ira Rubinstein and I am a Senior Fellow at the Information Law Institute, New York University School of Law, where I also teach Internet Law and Information Privacy Law as an Adjunct Professor. Over the past six months, I've been engaged in research on Privacy by Design (PbD). Specifically, I have been developing a framework for analyzing PbD, with the goal of developing recommendations for government policy makers. I therefore welcome the opportunity to submit comments in response to the Preliminary Staff Report of the Federal Trade Commission (FTC) on Protecting Consumer Privacy in an Era of Rapid Change: A Proposed Framework for Businesses and Policymakers (the "Staff Report").

The Staff Report describes a Proposed Framework with three components. In brief, it recommends that companies should 1) adopt PbD; (2) simplify consumer choice; and 3) increase the transparency of their data practices. PbD is the idea that "building in" privacy

throughout the design and development of products and services achieves better results than "bolting it on" as an afterthought. According the Staff Report (see pp. 44-52), companies engage in PbD when they promote consumer privacy throughout their organizations and at every stage of the development of their products and services. PbD has two main elements: 1) incorporating substantive privacy protections into a firm's practices such as data security, reasonable collection limitations, sound retention practices, and data accuracy; and 2) maintaining comprehensive data management procedures throughout the life cycle of their products and services. (These two elements correspond to a distinction discussed more fully below between "front-end" development practices and "back-end" data management practices.)

Staff specifically states that to ensure proper incorporation of the four substantive principles identified above, companies should develop and implement "comprehensive privacy programs," which in turn have two core elements: designating specific personnel with responsibility for privacy training and for promoting accountability for privacy policies; and conducting privacy impact assessments (PIAs), which companies may rely on to evaluate and mitigate risks. These privacy reviews should occur before a product launches and periodically thereafter to address any changes in data risks or other circumstances. The size and scope of comprehensive privacy programs should be determined by the risks presented to the data, with companies that collect vast amounts of consumer data or sensitive data required to devote more resources than those collecting small amounts of non-sensitive data. Finally, the report mentions in passing that staff supports the use of Privacy Enhancing Technologies (PETs) such as identity management, data tagging tools, transport encryption, and tools to "check and adjust default settings."

PbD as described in the Staff Report is certainly an enticing idea with great intuitive appeal. Why is this? The report suggests that by designing products and services with privacy in mind, companies are more likely to avoid or mitigate a security breach and to ensure that their products and services are better aligned with consumer expectations. There is indeed considerable evidence that resolving security issues during the design phase is more efficient and less costly than having to deal with it later in the development process.[1] It seems likely that this holds true of privacy issues as well. It also seems intuitively obvious that adding a new privacy feature late in the development process (or, even worse, after the product launches) increases costs just as it would to fix an error.

[1] *See* MARK GRAFF AND KENNETH VAN WYK, SECURE CODING: PRINCIPLES AND PRACTICES 56 (2003)(citing evidence that if the cost of a big fix at design time is taken as a unit of 1, the cost of fixing the same bug is about 6.5 times as great during implementation, 15 times as great during testing, and 60 times as great if a patch is required); s*ee generally* Nancy R. Mead, *Making the Business Case for Software Assurance* (2009)*, available at* https://buildsecurityin.us-cert.gov/bsi/articles/knowledge/business/685-BSI.html.

This point is readily confirmed by an example: Suppose that as a product nears the final release stage, a company decides that it had better add a new feature allowing customers to access their personal data, even though the product was not designed to expose personal data to end-users and lacks an appropriate and secure user interface. If the access feature requires additional or modified code or new user interfaces, or other additional functionality that was not planned for at the outset (such as authentication, which is needed to ensure that access is granted to the right person), then the costs of modification will increase even further.[2] At the very least, changing the program increases engineering time and effort beyond what would have been required if the feature had been designed in as the program was being written.

While I commend the Staff Report for endorsing PbD and agree with the recommendations described above, there are some gaps and oversights in the discussion that are worth noting before responding to the specific questions posed in the report. ***First,*** the report gives scant attention to PETs; ***second***, the report does not make it clear whether the main thrust of the PbD recommendation is that companies should perform PIAs or similar risk assessments anytime they release new products and services or significantly modify existing ones or should develop new rigorous privacy engineering practices including automated testing and compliance tools; and, ***third***, the topic of economic and regulatory incentives requires much fuller consideration. In addition to these gaps and oversights, the Staff Report takes insufficient advantage of the substantive analysis of privacy design issues already available from prior FTC enforcement actions, reports, and published guides, thereby missing an opportunity to offer its own preliminary set of best practices in privacy design and development. After briefly discussing these points, this letter concludes by responding to the questions on pages A1-A2.

1. What are PETs?

The report gives short shrift to PETs and—given their mixed record of success— perhaps this is appropriate. PETs have been around for about 25 years. Many PETs reflect major advances in cryptographic research that support advanced privacy features such as anonymous payment, various forms of anonymous (and security) protection for real-time communications, authentication via anonymous credential schemes, and methods for anonymously retrieving online content.[3] These PETs were first introduced as a regulatory strategy in 1995 in a joint

[2] For a useful (if somewhat dated) discussion of the complexities of designing online access features, *see* Final Report of the FTC Advisory Committee on Online Access and Security (May 15, 2000), *available at* http://www.ftc.gov/acoas/papers/finalreport.shtm.

[3] *See* Joan Feigenbaum et al, *Privacy Engineering in Digital Rights Management Systems,* ACM Workshop in Security and Privacy in Digital Rights Management (2001) *available at* http://cs-www.cs.yale.edu/homes/jf/FFSS.pdf .

report of the Dutch and Ontario Data Protection Authorities, appropriately entitled *Privacy-Enhancing Technologies: The Path to Anonymity*.[4] However, as a highly regarded group of computer scientists pointed out some fifteen years later, "Despite the apparent profusion of such technologies, few are in widespread use. Furthermore, even if they were in widespread use, they would not necessarily eliminate" user-privacy problems such as an overdependence on abstract models as opposed to "real-world" deployments; insecure implementations; ease-of-use issues; problems integrating PETs with legacy systems; and a variety of other user and technical issues. [5] Equally important are the economic aspects of anonymization tools: Because private firms profit from collecting customer data, they are more likely than not to reject any PET that would limit their access to this highly valuable information.

Of course, not all PETs rely on anonymity protocols. The term encompasses a range of tools beyond anonymity including those that enhance notice and choice (e.g., "just-in-time" notice, opt-in consent, and cookie managers), or help automate communication and/or enforcement of privacy policies via privacy languages (e.g., the Platform for Privacy Preferences (P3P) and the Enterprise Authorization Language (EPAL)). Some authors also classify as PETs various security tools such as those ensuring confidentiality via encryption.[6] Arguably, anonymity-based PETs are the most effective kind precisely because they prevent identification or collection of personal data irrespective of the requirements of privacy law. As a result, they are sometimes referred to as "true" or "pure" PETs. In contrast, many other PETs permit data collection and analysis but seek to enable knowledgeable and motivated consumers to exercise greater control over what data they share and with whom they share it.

Although the Commission recommends the use of PETs, the Staff Report needs to address PETs more fully and to clarify what it means when it says that it supports their use. The reasons for this are simple: First, a review of the relevant literature establishes that one of two basic meanings of PbD is simply incorporating the relevant PETs into a privacy design strategy.[7] Second, the costs of adopting PETs and their impact on online firms whose business models depend on the collection and analysis of data vary greatly depending on the nature and

---

[4] *See* Information and Privacy Commissioner, Ontario (Canada) and Registratiekamer (Netherlands), Privacy-Enhancing Technologies: The Path to Anonymity, Volume I (1995) *available at* http://www.ipc.on.ca/images/Resources/anoni-v2.pdf (hereinafter, "The Path to Anonymity").

[5] Feigenbaum, *supra* note 3 at 6.

[6] For recent overviews of the many varieties of PETs, *see* LONDON ECONOMICS, STUDY ON THE ECONOMIC BENEFITS OF PRIVACY-ENHANCING TECHNOLOGIES (PETS) 14-27 (2010) *available at* http://ec.europa.eu/justice/policies/privacy/docs/studies/final_report_pets_16_07_10_en.pdf; NATIONAL RESEARCH COUNCIL, ENGAGING PRIVACY AND INFORMATION TECHNOLOGY IN A DIGITAL AGE 107-16 (2007), *available at* http://books.nap.edu/catalog.php?record_id=11896.

[7] The other is a methodological approach, which defines PbD as an integrated set of development and management processes and practices; *see infra*, Section 3.

function of the relevant technology. In both cases, it matters a great deal what staff has in mind when they recommend PETs. In what follows, I suggest that PETs come in two very different breeds, which I will refer to as "substitute" and "complementary" PETs, and that the distinction between the two needs to be kept in mind in order to develop any useful guidance for the private sector.

2. Substitute vs. Complementary PETs

Substitute PETs seek to protect privacy by ensuring that little or no personal data is collected in the first place, thereby making legal protections superfluous. The main types of substitute PETs rely on anonymity to shield or reduce user identification and/or on client-centric architectures to prevent or minimize the collection of PII.[8] Their design is motivated by an underlying assumption that commercial IT systems are flawed, while legal rules and sanctions are in most (if not all) cases ineffective. Ideally, they offer privacy protections that make legal requirements irrelevant. Most of the best known substitute PETs are discrete applications deployed by individual end-users and provide limited functionality (e.g., anonymous browsing or encrypted email). Some substitute PETs also require ongoing maintenance, research and support from non-profits and volunteers (e.g., the Tor network).

In practice, many substitute PETs are more theoretical than practical and none are widely deployed for the reasons discussed above. Indeed, almost every firm that has sought to create a business around providing anonymity-based privacy tools or services has failed, which in turn discourages new firms from investing in them. Perhaps the most salient factor discouraging the broader deployment of these tools is economic self-interest: many of the most successful Internet firms have strong financial incentives to use customer data for targeted ads, personalization and/or differential pricing. Without the support of ISPs, e-commerce firms, search firms, and network advertising firms, all of which depend upon data collection and analysis for their core business, the commercial demand for substitute PETs is all but non-existent.

In sharp contrast, complementary PETs are designed to implement legislative privacy principles or related legal requirements. Thus, businesses are generally willing to deploy them both to ensure regulatory compliance and/or to give customers a positive impression of their commitment to privacy (defined as control over personal data). Developers of complementary PETs take it for granted that firms will collect data for various useful (and profitable) purposes but attempt to minimize potential consumer harms by ensuring that data is collected and

[8] *See* S. Spiekermann and L. Cranor, *Engineering Privacy,* 35 IEEE TRANSACTIONS ON SOFTWARE ENGINEERING 67 (2009).

processed in compliance with regulatory requirements based on FIPPs. Some complementary PETs focus on the "front-end" user experience (e.g., shorter, layered notices, informed and explicit consent, access and preference management tools), while others address privacy issues that arise with "back-end" infrastructure and data sharing networks (e.g., IBM's Tivoli Privacy Manager, which helps enterprises manage user identities, access rights and privacy policies across an entire e-business infrastructure, and HP's proposed Policy Compliance Checking System).

Complementary PETs may be classified into two sub-categories: First, *privacy-friendly* PETs, whose overall goal is to give people more control over their personal data through improved notice and consent mechanisms, browser management tools, marketing preference dashboards, and so on; and second, *privacy-preserving* PETs, which (in many cases) resemble substitute PETs in relying on sophisticated cryptographic protocols that may lead to deployable solutions with strong privacy guarantees but that also complement legal requirements. This combination of features permits companies (and government agencies) to engage in activities that might otherwise be viewed as privacy invasive while preserving privacy in a rigorous manner. Good examples include privacy-preserving data mining,[9] privacy-preserving *targeted* advertising,[10] and differential privacy.[11]

The distinction between substitute and complementary PETs is nicely illustrated by PETs designed to control the receipt of targeted advertising (*see* Staff Report, pp. 63-69). Examples of relevant tools in each of the main categories of PETs identified above include the following:

1. Substitute PETs: Various anonymity tools are available that would prevent tracking and targeted advertising by enabling consumers to surf the web anonymously. For example, anonymous proxy servers permit users to surf the web without revealing their IP addresses; the Tor Browsing bundle offers similar functionality using a much stronger cryptographic protocol. Consistent with their business models, however, none of the major search or network advertising firms support the use of such tools, directly or indirectly, in their web services.

[9] *See, e.g.,* R. Agrawal and R. Srikant, *Privacy-Preserving Data Mining,* 29 SIGMOD RECORD 439 (2000), *available at* http://www.cs.utexas.edu/~shmat/courses/cs395t_fall05/ppdm.pdf

[10] V. Toubiana et al *Adnostic: Privacy preserving targeted advertising*, in 17th Annual Network and Distributed System Security Symposium, NDSS, 2010, *available at* http://crypto.stanford.edu/adnostic/adnostic.pdf.

[11] *See, e.g.,* P. Golle et al, *Data Collection with Self-Enforcing Privacy*, Proc. 13th ACM Conf. Computer and Comm. Security, pp. 69-78 (2006), *available at* http://research.microsoft.com/en-us/projects/databaseprivacy/self-full.pdf.

2. Complementary (Privacy-Friendly) PETs: On the other hand, many of the most popular commercial Internet and network advertising firms strongly support tools that enable users to control their online advertising by editing their inferred interest and demographic categories or opting-out of behavioral targeting with respect to participating firms. Examples include ad preference managers from Google and eXelate, standalone and browser-based cookie managers, additional browser controls that allow users to delete cookies (including Flash cookies), "private browsing" features (which delete cookies each time the user closes the browser or turns off private browsing, effectively hiding his or her history), new icons that link to additional information and choices about behavioral advertising, and new, browser-based "do not track" tools from all three of the major browser vendors.

3. Complementary (Privacy-Preserving) PETs: Finally, a group of researchers recently developed a privacy-preserving approach to targeted advertising, which they call Adnostic.[12] This proposed system would allow ad networks to engage in behavioral profiling and ad targeting but without having a server track consumers. Rather, all of the tracking and profiling necessary for serving targeted ads takes place on the client-side, i.e., in the user's own browser. When a site wants to serve an interest-based ad, the user's browser chooses the most relevant ad from a portfolio of ads offered by the ad network service but the browser doesn't reveal this information to the ad service (or to any third-party).[13]

Why are these distinctions important? The answer relates to the incentives for developing and using PETs. As already noted, firms that have financial incentives to collect and analyze consumer data may be extremely reluctant to adopt substitute PETs. Indeed, the market incentives for such PETs are very limited or even non-existent and it seems unlikely that FTC could persuade businesses to deploy them unless it threatened and had the authority to promulgate highly restrictive regulations. On the other hand, a much stronger business case exists for (privacy-friendly) complementary PETs because they both support existing compliance obligations and tend to enhance a firm's reputation as a trustworthy company that cares about privacy. Indeed, these PETs are attractive to companies for obvious reasons: they enhance notice and choice in a privacy-friendly manner but without disrupting the advertising business model. Adnostic is also very promising precisely because it offers much greater privacy protections than privacy-friendly PETs without undermining the advertising business models.

---

[12] *See supra* note 10.

[13] *Compare* Ann Cavoukian, *Redesigning IP Geolocation:* Privacy by Design *and Online Targeted Advertising* (2010), *available at* *http://www.ipc.on.ca/english/Resources/Discussion-Papers/Discussion-Papers-Summary/?id=985* (discussing Bering Media's "double-blind" privacy architecture).

On the other hand, Adnostic imposes new costs and complexity on the online advertising industry and arguably undermines the ability of different ad services to compete based on which of them has the best matching algorithms, which may limit its adoption. Of course, business will adopt complementary PETs only if they determine that the (direct and opportunity) costs of doing so are low enough to justify the investment. But there is evidence that firms will under-invest in privacy technology if they can do so without it hurting their bottom line. Thus, regulatory incentives may still be necessary to overcome the reluctance of private firms to increase their investments in privacy technology in the face of competing business demands and a weak economy.[14] This seems especially true if the Commission wishes to encourage businesses to adopt privacy-preserving, as opposed to privacy-friendly, PETs.

3. Two Approaches to PbD

What does it mean for a firm to build privacy into its products and services? The Staff Report does not offer a precise definition but based on staff's discussion, "building in" privacy seems to refer to 1) a prospective activity in which systems or application designers 2) think through the privacy implications of their work in a systematic fashion, 3) pay attention to privacy throughout the product development and data management lifecycles, and 4) rely on risk assessments to determine the appropriate level of privacy protections. But these are very general requirements and without more detailed guidance, firms interested in pursuing a PbD strategy will not know what they are supposed to do (and not do).

Arguably, PbD may be thought of in two very different ways. First, as noted above, PbD may be viewed simply as the use of existing PETs or the creation of new PETs in response to new privacy needs. If this is what the Commission has in mind, the Staff Report needs to identify those PETs that would help firms implement FIPPs-based privacy protections. In addition, the report might recommend PETs for use under specific circumstances. But this type of analysis and evaluation is largely missing from the report, which discusses privacy-friendly choice mechanisms for online behavioral advertising (including "Do Not Track") but otherwise fails to identify or discuss more broadly any substitution or privacy-preserving PETs.

Second, PbD may be defined in methodological terms. On this view, PbD consists in an integrated set of development and management processes and practices. As with PETs, it is necessary to differentiate front-end software development activities from back-end data management practices. The *software development lifecycle* seeks to ensure that in designing products and services, software developers take account of both customer privacy expectations and the relevant threat model that needs to be guarded against. This approach empowers users

[14] Economic and regulatory incentives are briefly discussed below, *infra*, Section 4.

to control their personal data (for example, by improving their understanding of what information will be collected from them, how it will be used and what choices they have as to its transfer, storage and use) and seeks to minimize the risks of privacy incidents (such as surreptitious or unanticipated data collection, unauthorized data use, transfer or exposure, or security breaches). The *data management lifecycle*, on the other hand, focuses more on how to engineer and manage information systems with privacy in mind as a company's own employees and business partners access, use, disclose, and eventually delete customer data in the ordinary course of operating a business. The former is a design process for customer-facing products and services (i.e., those with which customers interact by downloading software, using a web service, and/or sharing personal data or creating user content); the latter consists in data management processes and practices that ensure that information systems (for both internal use and for sharing data with affiliates, partners, and suppliers) comply with privacy laws, company policies (including published privacy policies), and customers' own privacy preferences. Although distinctive, the two lifecycles overlap given that most products and services designed for the Internet also depend on back-end data handling.

This front-end/back-end distinction is generally consistent with the chief concerns discussed in sections V(B)(1) and V(B)(2) of the Staff Report. The former advises companies on "incorporating substantive privacy practices into their practices," while the latter recommends that companies maintain "comprehensive data management procedures." Yet there are problems with both sections. The main shortcomings are a lack of detail or actionable guidance, which in turn makes it unclear whether the Commission's main recommendation is that business adopt PIAs or invest in privacy engineering including automated testing and compliance tools. To see this, it is instructive to compare these sections with corresponding guidelines and policies developed in the private sector.

Several of the older and more well-established, multinational IT companies have developed guidelines, policies, tools, and systems for building privacy into software development and data management. For example, Microsoft's "Security Development Lifecycle" (SDL) for software development is the best known example of how privacy can be built into the design process. The SDL aims to integrate privacy and security principles into the software development lifecycle, but each of the five stages of the development lifecycle (requirements, design, implementation, verification, and release) also includes privacy guidelines, which range from the mandatory to the recommend and from the procedural to the technical. Privacy impact ratings are given to each project and these ratings determine the design specifications needed for compliance. The SDL guidelines are supplemented by Microsoft's "Privacy Guidelines for Developing Software and Services," a 51-page document that lays out basic concepts and definitions based on the FIPPs and related US privacy laws; discusses different types of privacy

controls and special considerations raised by shared computers, third parties, and other situations; and then enumerates detailed guidelines for nine specific software product and Web site development scenarios.[15] For each scenario, the guidelines identify required and recommended practices relevant to notice and consent, security and data integrity, customer access, use of cookies, and additional controls or requirements.[16]

On the data management side, IBM's Tivoli Privacy Manager is a comprehensive enterprise privacy management system that supports a variety of privacy functionalities.[17] HP is also developing a comprehensive approach to managing the information lifecycle (storage, retrieval, usage, prioritization, update, transformation, and deletion) as well as identity management tasks (such as the collection, storage, and processing of identity and profiling information, authentication and authorization, "provisioning" of digital identities (i.e., account registration and related tasks), and user management of personal data and identities). According to researchers in HP's Trusted Systems Lab, this requires both a model of privacy obligations (based on the rights of data subjects, any permission they have granted over the use of their personal data, and various statutory obligations associated with the FIPPs) and a framework for managing these obligations. The resulting "obligation management system" enables enterprises to configure information lifecycle and identity management solutions to deal with the preferences and constraints dictated by privacy obligations and ideally to do so in an automated and integrated fashion.[18]

---

[15] *See* Microsoft Privacy Guidelines for Developing Software Products and Services, v. 3.1 (Feb. 2008), *available at* http://www.microsoft.com/downloads/en/details.aspx?FamilyID=c48cf80f-6e87-48f5-83ec-a18d1ad2fc1f&DisplayLang=en (describing nine scenarios at length including Transferring PII to and from, Storing PII on, Transferring Anonymous Data from, and Installing Software on a Customer's System, Deploying a Web Site, Storing and Processing User Data at, and Transferring User Data outside, the Company, Interacting with Children, and Server Deployment). *See also* Tina R. Knutson, "Building Privacy into Software Products and Services," *IEEE Security and Privacy*, vol. 5, no. 3, pp. 72-74, May/June 2007.

[16] Although the guidelines mainly treat privacy design issues for front-end products and services, they also address such back-end services as "Server Deployment" (Scenario 9). This reinforces my earlier suggestion that the front-end/back-end distinction is not exclusive but instead describes primary areas of focus.

[17] *See* Paul Ashley and David Moore, Enforcing Privacy Within an Enterprise Using IBM Tivoli Privacy Manager for E-business (2002), *available at* http://www.ibm.com/developerworks/tivoli/library/t-privacy/index.html (describing functions such as tracking different versions of privacy policies; storing consent of the individual to the privacy policy when PII data is collected; auditing of all submissions and accesses to PII; and authorization of submissions and accesses to PII). *See also* Roberto J. Bayardo and Ramakrishnan Srikant, *Technological Solutions for Protecting Privacy,* 36 COMPUTER 115 (2003), *available at* http://www.almaden.ibm.com/cs/projects/iis/hdb/Publications/papers/ieee03.pdf.

[18] *See* Marco Mont, *On Privacy–Aware Information Lifecycle Management in Enterprises: Setting the Context*, HPL-2006-109 (2006), *available at* http://www.hpl.hp.com/techreports/2006/HPL-2006-109.html (describing the five core properties and functionalities of a privacy–aware information lifecycle management solution as including explicit modeling of personal and confidential data; explicit definition of privacy policies, in particular obligations; integrated lifecycle management of these policies; deployment and enforcement of these policies, potentially by leveraging information lifecycle management and identity management infrastructures; and integrated monitoring and checking for compliance to these policies).

In comparison to these commercial approaches, which are both rich in detail and very comprehensive, the discussion of privacy development guidelines in Section V(B)(1) of the Staff Report seems incomplete. To begin with, it considers only four substantive privacy protections that firms should incorporate into their practices (security, collection limits, retention practices and accuracy) but fails to explain why all eight FIPPs are not applicable.[19] Certainly, two of these principles—purpose specification and use limitation—are highly relevant to building privacy protections into products and services. An equally serious omission of this section (but not of later sections of the report) is the failure to discuss common use scenarios or the rules that should govern them, the severity of threat associated with each of them, and the safeguards needed to address these threats consistent with customer expectations and legal requirements.[20]

In Section V(B)(2), the report's guidance consists mainly in recommending first, that firms implement "comprehensive privacy programs" and, second, that they assess risks (in a manner akin to PIAs) "where appropriate." But these insights are not sufficiently developed to provide any useful guidance. For example, the report neglects to define when risk assessments *are* appropriate. This is surprising considering that Section 208(b)(1)(A) of the E-Government Act of 2002 offers relevant guidelines, requiring federal agencies to perform a PIA prior to developing or procuring IT systems or projects that collect, maintain or disseminate information in identifiable form from or about members of the public.[21] Although the Staff Report offers a few illustrations of privacy reviews (notably in its discussion of peer-to-peer file sharing), and a little

[19] *See, e.g., OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data,* O.E.C.D. Doc. (C 58 final) (Oct. 1, 1980) (identifying eight basic principles as follows: Collection Limitation, Data Quality, Purpose Specification, Use Limitation, Security Safeguards, Openness, Individual Participation, and Accountability; U.S. Department of Homeland Security, *Privacy Policy Guidance Memorandum, The Fair Information Practice Principles: Framework for Privacy Policy at the Department of Homeland Security* (Dec. 2008)(identifying a very similar set of eight principles).

[20] In fact, Sections V(C) and (D) of the Staff Report examine a number of scenarios involving choice, notice, access and material changes (see pp. 58-77) but the report does not refer to this analysis in its discussion of PbD.

[21] *See* OMB, Guidance for Implementing the Privacy Provisions of the E-Government Act of 2002, M-03-22 (Sep. 26, 2003)(further specifying that PIAs are required "to be performed and updated as necessary where a system change creates new privacy risks" such as conversions of paper-based records to electronic systems; when functions applied to an existing information collection change anonymous information into information in identifiable form; when there are significant system management changes; when there is significant merging of government databases; when there are significant new uses or exchanges of information in identifiable form; when alteration of a business process results in significant new uses or disclosures of information or incorporation into the system of additional items of information in identifiable form; and when new information in identifiable form added to a collection raises the risks to personal privacy (for example, the addition of health or financial information).

(very little) prescriptive guidance,[22] it does not go far enough in providing detailed rules or requirements for privacy assessments to help companies determine when to conduct assessments or whether they have done so in a meaningful way.

The report also fails to describe the proposed "comprehensive privacy programs" with the same level of detail as the analogous "information security programs" as defined in both the Safeguards Rule of the Gramm-Leach-Bliely Act[23] and numerous FTC enforcement actions (which are cited by staff at pp. 10-11). For example, a recent FTC Consent Order requires Twitter to "establish and maintain a comprehensive information security program in writing that is reasonably designed to protect the security, privacy, confidentiality, and integrity of nonpublic consumer information. The security program must contain administrative, technical, and physical safeguards appropriate to Twitter's size and complexity, the nature and scope of its activities, and the sensitivity of the nonpublic consumer information."[24] The order specifically requires Twitter to meet five requirements:

1. Designation of a responsible employee to coordinate and be accountable for the information security program;
2. Risk assessment (defined as identifying reasonably-foreseeable risks that could result in "the unauthorized disclosure, misuse, loss, alteration, destruction, or other compromise of nonpublic consumer information or in unauthorized administrative control of the Twitter system";
3. Design and implementation of reasonable safeguards to control the risks identified through risk assessment, and regular testing or monitoring of the effectiveness of the safeguards' key controls, systems, and procedures;
4. Reasonable care in selecting and retaining service providers; and
5. Evaluation and adjustment of the program as a result of testing or any material changes.

At a high level of generality, there is considerable overlap between these security requirements (especially requirements 1, 2 and 5 above) and the elements of a comprehensive privacy program. Perhaps the Commission's reluctance to provide more detailed guidance regarding comprehensive privacy programs stems from its lack of regulatory authority as compared to

[22] For example, "companies dealing with consumers' data should keep up-to-date on privacy-related developments and should modify their practices as necessary to maintain privacy and ensure that their practices comport with their representations to consumers" (p. 52).

[23] The Safeguards Rule, 16 C.F.R. pt. 314, implements the security and confidentiality requirements of the Gramm-Leach-Bliely Act, 15 U.S.C. §§ 6801-09, as applied to "nonpublic personal information" or NPI. This Rule requires that information security programs seek to (1) insure the security and confidentiality of NPI; (2) protect against any anticipated threats or hazards to the security or integrity of NPI; and (3) protect against unauthorized access to or use of NPI that could result in substantial harm or inconvenience to any customer; *see* 16 C.F.R. § 314.3 (2008).

[24] *See* Agreement Containing Consent order, *In re* Twitter, Inc., File No. 0923093 (June 24, 2010).

GLB, which grants explicit rulemaking authority to the FTC regarding financial institutions safeguards.[25] As discussed below in greater detail, however, the Commission has other options for issuing guidance even absent this explicit authority. For example, it might develop guidelines based on its own privacy enforcement cases alleging that specific software practices are unfair or deceptive practices under Section 5 of the FTC Act. Or it might establish an advisory committee on privacy by design, whose charter would include providing advice and recommendations to the Commission regarding implementation of privacy design principles and practices by domestic commercial firms that collect and analyze data. This committee might then evaluate and recommend not only on the elements of a privacy program generally but specific privacy design practices that might be classified as prohibited, required or recommended.[26]

Finally, this comparison between comprehensive privacy and information security programs raises an important point that is not discussed in the Staff report. The security requirements identified in the Safeguards Rule and FTC enforcement cases are not only more detailed than anything discussed in the Staff Report but identify specific risks to guard against (unauthorized disclosure, misuse, loss, alteration, destruction, or other compromise of information) and they include a testing requirement. Security engineers understand what it means to take account of such risks as they design products and services. They rely not only on risk assessments but on well-accepted security design principles[27] and security techniques including threat modeling, secure coding, and secure testing practices (such as fuzz testing, penetration testing, and run-time verification).[28] In contrast, privacy design principles are less well-established and privacy engineering techniques are still in their infancy. This lack of tried-and-true principles and techniques for achieving privacy by design imply an even greater need on the part of FTC to provide more extensive guidance on what it has in mind. In particular, in recommending PbD as a part of its Proposed Framework, staff needs to clarify whether it is suggesting that the private sector should implement PIAs as have federal government agencies (with mixed results),[29] or if

[25] *See* 15 U.S.C. §§ 6801(b) and 6804(a)(1).

[26] *See infra* Appendix A.

[27] For a classic statement of these principles, *see* Jerome H. Saltzer and Michael D. Schroeder, *The Protection of Information in Computer Systems*, Proceedings of the IEEE 63, 9 (September 1975), pp. 1278-1308

[28] S*ee, e.g.,* MICHAEL HOWARD AND DAVID LEBLANC, WRITING SECURE CODE (2000); MARK GRAFF AND KENNETH VAN WYK, SECURE CODING: PRINCIPLES AND PRACTICES (2003); MICHAEL HOWARD AND STEVE LIPNER, THE SECURITY DEVELOPMENT LIFECYCLE (2006); GARY MCGRAW, SOFTWARE SECURITY: BUILDING SECURITY IN (2006).

[29] *See* Roger Clarke, *An Evaluation of Privacy Impact Assessment Guidance Documents* (2010), *available at* http://www.rogerclarke.com/DV/PIAG-Eval.html#App (noting that while PIAs are mandated under the E-Government Act, the statute and the OMB Guidance "limit the scope of the assessment to legal compliance, no consultation is required (nor does it appear that any is ever undertaken), and although reports are publicly available they are checklist-based and almost entirely devoid of any content of significance to privacy protection,

they should develop and follow privacy development guidelines and implement data lifecycle management practices and automated testing and compliance solutions as have major firms such as Microsoft, IBM and HP. The difference has to do with whether PbD is ultimately a risk assessment activity conducted by privacy experts or an engineering discipline carried out by software developers in consultation with privacy officers, lawyers and other with expertise in regulatory compliance.

4. Market and Regulatory Incentives

As noted above, firms have strong market incentives to collect and analyze customer data for profit-making activities such as targeted ads, personalization and differential pricing. This severely limits the adoption of certain PETs or of PbD methods that might restrict access or use of customer data. Nor is it obvious how firms should determine whether they would benefit from investing in privacy technology. This is mainly due to the fact that a cost-benefit analysis requires relevant data such as the cost of privacy breaches and probability of their occurrence. But there are no standardized measures of these costs and very few surveys of the costs of privacy.[30] In the security field, researchers have shown that in the absence of cost data, firms are less likely to undertake a cost-benefit analysis at all. Instead, they are more likely to adjust security budgets up or down as dictated by management, often for unrelated reasons, or to adopt a passive strategy, in which they increase security technology expenditures only in response to a data breach or major security incident or anticipated lost sales associated with strongly negative customer perceptions of their commitment to security.[31] It seems reasonable to assume that firms behave in a similar manner in deciding on privacy investments.

The high value that firms place on data collection and use, together with the obstacles to applying cost-benefit thinking to privacy investments, perhaps explains why firms generally shun substitute PETs and seem to prefer privacy-friendly PETs over privacy-preserving PETs.[32] Moreover, firms may delay investments in the processes and practices that would help build privacy into front-end products and services or back-end data management systems, either

---

beyond the narrowly circumscribed legal requirements. Where agencies have their own internal guidance documents, they are also uniformly limited to compliance checks").

[30] One of the very few such studies is IBM AND PONEMON INST., THE COSTS OF PRIVACY STUDY (Feb. 17, 2004).

[31] *See, e.g.,* Lawrence Gordon and Martin Loeb, *Budgeting Process for Information Security Expenditures,* 49 COMMUNICATIONS OF THE ACM 121 (June 2006).

[32] This is only a general impression and more research is required to substantiate it.

because, in the absence of data supporting a cost-benefit analysis, it is difficult to justify these expenditures, or because there is a perceived lack of customer interest in privacy.[33]

Even in the face of weak economic incentives for widespread adoption of privacy technology, there is a road map for FTC to follow in creating more powerful regulatory incentives for firms to implement PbD. This road map consists in the Commission playing the same activist role in developing a comprehensive privacy program as it did in spelling out the elements of a comprehensive information security program. As suggested below, the Commission may rely on its existing authority under Section 5 of the FTC Act to 1) bring enforcement cases against businesses alleging a failure to take reasonable steps to incorporate substantive privacy protections into their design practices and development lifecycle or to maintain comprehensive data management procedures throughout the data lifecycle and 2) to convene an advisory committee and/or workshops and roundtables for the purpose of identifying best practices in PbD, and issuing appropriate guidance to businesses based on this learning.

Alternatively, if Congress enacts a new privacy law making FIPPs broadly applicable to firms that collect PII and requiring firms to implement FIPPs into their design practices and data management procedures, and authorizing the FTC to establish a "co-regulatory" safe harbor program, that is, a program that incentivizes organizations to meet high standards of data protection by shielding safe harbor participants from various "sticks" (such as a private right of action), and rewarding them with various "carrots" (such as by allowing greater flexibility in how they implement certain requirements), then the Commission would be well-positioned to further develop and specify the requirements of a comprehensive privacy program by regulation.[34]

[33] Although polls suggest growing consumer concerns over privacy, firms may justify their low investment in PETs and PbD by arguing that consumers' actual behavior is quite different: they do little to protect their own privacy and may sacrifice it in exchange for small rewards. For an alternative explanation of this divergence between revealed preferences and actual behavior, *see* Alessandro Acquisti and Jens Grossklags, *Privacy and Rationality in Individual Decision Making*, 3 IEEE Security & Privacy 26 (2005)(discussing factors such as incomplete information, bounded rationality and behavioral biases). The cost-benefit trade-offs or disclosing or protecting customer data is a complex topic and beyond the scope of these comments; for a more comprehensive discussion, *see* Alessandro Acquisti, *The Economics of Personal Data and the Economics of Privacy* (2010), *available at* http://www.oecd.org/dataoecd/8/51/46968784.pdf .

[34] For an example of a privacy bill requiring that firms implement FIPPs and providing a safe harbor option, *see* the BEST PRACTICES Act, H.R. 5777, 111th Cong. (2010). For a broader discussion of "co-regulatory" safe harbors and what they might contribute to the privacy debate, *see generally* Ira Rubinstein, *Privacy and Regulatory Innovation: Moving Beyond Voluntary Codes*, NYU School of Law, Public Law Research Paper No. 10-16 (March 1, 2010), *available at* http://ssrn.com/abstract=1510275.

5. Are Best Practices the (Provisional) Answer?

The preceding analysis suggests that staff has more work to do to flesh out the PbD recommendation currently contained in the Staff Report. Unfortunately, in the absence of tried-and-true privacy design principles, this is not an easy task. Assuming that Congress does not enact new privacy legislation requiring firms to adopt PbD and authorizing the Commission to issue appropriate regulations, either under a safe harbor model or otherwise, the next best step would be for FTC to bring together industry, academic and advocates with the necessary expertise in an advisory committee,[35] or at a workshop designed to identify and debate best practices. Through an advisory committee report, workshops, round tables, and other guidance, the Commission could then begin to fill in the details currently missing from the Staff Report. The Commission might also consider supporting ongoing efforts by the ISO and others to define international privacy design standards.

Another obvious source of guidance is the Commission's own privacy enforcement actions, especially cases alleging that specific software practices are unfair under Section 5 of the FTC Act. Indeed, the Commission's published decisions (and no-action letters) already offer some guidance on prohibited, required and recommended privacy design practices. Interestingly, the Staff Report mentions the almost thirty cases it brought against companies for alleged failures of security practices. The orders obtained in these cases require companies to implement comprehensive information security programs and to obtain third-party audits of the effectiveness of those programs" (citations omitted).[36] But while the Staff Report discusses a recent action against Sears involving the privacy implications of tracking software, it does not discuss several other privacy and spyware cases alleging this and other unfair software practices which, taken together, and in combination with staff's discussion of "commonly accepted" practices in providing notice and choice (see pp. 53-63), and how to increase transparency in data practices (pp. 69-77), begin to suggest the outlines of a set of best practices in privacy design. By developing best practices based on both existing cases and other materials, and as developed in future enforcement actions and FTC-sponsored workshops, round tables and other presentations, the Commission might flesh out its recommended

[35] In December 1999 the Federal Trade Commission established the Advisory Committee on Online Access and Security. The purpose of the Advisory Committee was to provide advice and recommendations to the Commission regarding implementation of certain fair information practices by domestic commercial Web sites. In particular, the Advisory Committee addressed providing online consumers reasonable access to personal information collected from and about them and maintaining adequate security for that information. The Federal Register Notice establishing the committee, the committee charter, and other relevant documents are *available at* http://www.ftc.gov/acoas/index.shtm.

[36] Staff Report, 10-11.

"comprehensive privacy program," thereby making it as detailed and actionable as the analogous "information security program."[37]

Conclusion

This comment concludes with a set of brief responses to those specific questions identified in the Staff Report regarding PbD and which are discussed above. (See Appendix B.)

Please do not hesitate to contact me if I can be of any further assistance in developing and elaborating the recommendation that firms implement privacy by design.

Sincerely,

Ira S. Rubinstein
Senior Fellow, Information Law Institute
New York University School of Law

[37] Appendix A represents a first effort at collecting relevant enforcement cases that bear on software design practices and classifying them under the headings of prohibited, required or recommended practices.

Appendix A: Relevant Enforcement Cases

1. *Prohibited practices.* Companies shall not:
    a. Exploit any security vulnerability to download or install software;[38]
    b. Distribute software code bundled with "lureware" that tracks consumers' Internet activity or collects other personal information, changes their preferred homepage or other browser settings, inserts new toolbars onto their browsers, installs dialer programs, inserts advertising hyperlinks into third-party Web pages, or installs other advertising software;[39]
    c. Install content protection software that that hides, cloaks or misnames files, folders, or directories, or misrepresents the purpose or effect of files, directory folders, formats, or registry entries.[40]

2. *Required practices.* Companies must:
    a. Clearly and conspicuously disclose when free software is bundled with harmful software (malware) creating security and privacy risks for consumers who install it;[41]
    b. Clearly and conspicuously disclose that the installation of software from a CD may limit a consumer's ability to copy or distribute audio files from the CD or other digital content; and, if such software causes information about consumers, their computes, or their use of a product to be transmitted via the Internet (so-called "phone home" features), then companies must disclose this prior to any such transmission and obtain the consumer's opt-in consent;[42]
    c. Provide a readily identifiable means for consumers to uninstall any adware or similar programs that monitor consumers' Internet use and display frequent, targeted pop-up ads, where the companies deliberately made these adware programs difficult for consumers to identify, locate, and remove from their computers.[43]
    d. Clearly and prominently disclose the types of data that certain tracking software will monitor, record, or transmit prior to installing this software and separate from any user license agreement. Sears also must disclose whether any data will be used by a third party.[44]

[38] *FTC v. Odysseus Marketing, Inc.*, No. 05-CV-330 (D.N.H. Oct. 24, 2006).
[39] *FTC v. Enternet Media, Inc.*, CV 05-7777 CAS (C.D. Cal., Aug. 22, 2006).
[40] Decision and Order, *In re* Sony BMG Music Entm't, FTC Docket No. C-4195 (June 28, 2007).
[41] Consent Order, *In re* Advertising.com, FTC File No. 042 3196 (Sept. 12, 2005).
[42] Decision and Order, *In re* Sony BMG Music Entm't, , FTC Docket No. C-4195 (June 28, 2007).
[43] Decision and Order, *In re* Zango, , FTC Docket No. C-4186 (March 9, 2007).
[44] *In re* Sears Holdings Management Corporation, FTC File No. 082 3099 (Sept. 9, 2009).

e. Provide prominent disclosures and obtain opt-in consent before using consumer data in a materially different manner than claimed when the data was collected, posted, or otherwise obtained.[45]

3. *Recommended practices.* Companies should:
    a. Develop and implement reasonable procedures concerning the collection and use of any personally identifiable information, including collecting information only to the extent necessary to fulfill a business purpose, disposing of the information no longer necessary to accomplish that purpose, and maintaining the privacy and security of information collected and stored.[46]
    b. Incorporate a formal privacy review process into the design phases of new Initiatives;[47]
    c. Where a company has a relationship with a consumer, it should offer a choice mechanism "at the point when the consumer is providing data or otherwise engaging with the company" (Staff Report, p. 58).
    d. Where a social media firm conveys consumer information to a third-party application developer, "the notice-and-choice mechanism should appear at the time the consumer is deciding whether to use the application and in any event, before the application obtains the consumer's information" (Staff Report, p. 59).
    e. Where consumers elect not to have their information collected, used, or shared, "that decision should be durable and not subject to repeated additional requests from the particular merchant" (Id.).
    f. Seek affirmative express consent before collecting, using, or sharing any "sensitive information" including "information about children, financial and medical information, and precise geolocation data" (Staff Report, p. 61).
    g. Where companies are engaged in online behavioral advertising, they should use a special choice mechanism consisting in "Do Not Track" (Staff Report, 63-69).
    h. "Privacy notices should provide clear, comparable, and concise descriptions of a company's overall data practices" (Staff Report, 71).
    i. Implement a "sliding scale" approach to access, taking into account the costs and benefits of access in different situations (Staff Report, 72-73).

[45] *See Gateway Learning Corp.*, No. C-4120, 2004 WL 2618647 (F.T.C. Sept. 10, 2004)

[46] Letter to Albert Gidari, Esq., Counsel for Google, From David C. Vladeck, Director, Bureau of Consumer Protection, Closing Google Inquiry (Oct. 27, 2010).

[47] *Id.*

Appendix B: Questions for Comment (PbD)

Companies should promote consumer privacy throughout their organizations and at every stage of the development of their products and services

Incorporate substantive privacy protections

• Are there substantive protections, in addition to those set forth in Section V(B)(1) of the report, that companies should provide and how should the costs and benefits of such protections be balanced?

Yes. Companies should provide all of the protections required by the OECD Guidelines, as expressed in the FIPPs adopted by DHS. At a minimum, these protections should include purpose specification and use limitations principles.

As to the balancing of costs and benefits, at a minimum, a reasonableness standard should apply as is does in the case of information security programs. More generally, the costs and benefits of comprehensive privacy programs are not well-known. Recently, a consulting firm called London Economics completed a Final Report to the European Commission, DG Justice, Freedom and Security, on the economic benefits of privacy-enhancing technologies (PETs).[48] *More surveys and empirical investigations of this nature are needed before costs and benefits of PbD can be assessed.*

• Should the concept of "specific business purpose" or "need" be defined further and, if so, how?

"Specific business purpose" is a broad context and covers a multiplicity of specific contexts, which makes it difficult to define. Moreover, any purpose specification needs to be read in combination with both data minimization and use limitation principles. These principles, as formulated by the Department of Homeland Security (DHS) in 2008,[49] *might be recast as guidance directed at the private sector as follows:*

[48] Available at http://ec.europa.eu/justice/policies/privacy/docs/studies/final_report_pets_16_07_10_en.pdf.

[49] *See* U.S. Department of Homeland Security, *Privacy Policy Guidance Memorandum, The Fair Information Practice Principles: Framework for Privacy Policy at the Department of Homeland Security* (Dec. 2008), *available at* http://www.dhs.gov/xlibrary/assets/privacy/privacy_policyguide_2008-01.pdf ("DHS FIPs").

Purpose Specification. *Companies should specifically articulate the purpose or purposes for which personal information is intended to be used.*
Data Minimization. *Only data directly relevant and necessary to accomplish a specified purpose should be collected and data should only be retained for as long as is necessary to fulfill a specified purpose.*
Use Limitation. *Personal information should be used solely for the purpose(s) specified in the notice. Sharing of personal information should be for a purpose compatible with the purpose for which it was collected.*

• Is there a way to prescribe a reasonable retention period?
[No opinion as this issue was not discussed above.]

• Should the retention period depend upon the type or the sensitivity of the data at issue? For example, does the value of information used for behavioral advertising decrease so quickly that retention periods for such data can be quite short?
[No opinion as this issue was not discussed above.]

• How should the substantive principles set forth in Section V(B)(1) of the report apply to companies with legacy data systems?
Legacy systems are often difficult and costly to update if new functionality is required and this applies to privacy requirements as well. Often, firms need a catalyst to replace legacy systems, which might consist in new market requirements, new competitors, or new legal requirements. Thus, the FTC should design appropriate regulatory incentives that , on the one hand, encourage firms to update legacy systems in response to market pressures and, on the other hand, establish deadlines for upgrading or replacing systems that fail to implement substantive privacy protections.

• When it is not feasible to update legacy data systems, what administrative or technical procedures should companies follow to mitigate the risks posed by such systems?
[This is a technical issue and beyond the scope of these comments.]
• Can companies minimize or otherwise modify the data maintained in legacy data systems to protect consumer privacy interests?
[This is a technical issue and beyond the scope of these comments.]

Maintain comprehensive data management procedures

• How can the full range of stakeholders be given an incentive to develop and deploy privacy-enhancing technologies?

As discussed above (supra, pp. 6-7), this question can only be answered with respect to specific categories of PETs.

• What roles should different industry participants – *e.g.*, browser vendors, website operators, advertising companies – play in addressing privacy concerns with more effective technologies for consumer control?
As discussed above (supra, pp. 6-7), this question can only be answered with respect to specific categories of PETs.