Sci Eng Ethics
DOI 10.1007/s11948-015-9674-9

ORIGINAL PAPER

Respecting Context to Protect Privacy: Why Meaning Matters

Helen Nissenbaum[1]

Received: 7 March 2015 / Accepted: 22 June 2015

Abstract In February 2012, the Obama White House endorsed a Privacy Bill of Rights, comprising seven principles. The third, "Respect for Context," is explained as the expectation that "companies will collect, use, and disclose personal data in ways that are consistent with the context in which consumers provide the data." One can anticipate the contested interpretations of this principle as parties representing diverse interests vie to make theirs the authoritative one. In the paper I will discuss three possibilities and explain why each does not take us far beyond the status quo, which, regulators in the United States, Europe, and beyond have found problematic. I will argue that contextual integrity offers the best way forward for protecting privacy in a world where information increasingly mediates our significant activities and relationships. Although an important goal is to influence policy, this paper aims less to stipulate explicit rules than to present an underlying justificatory, or normative rationale. Along the way, it will review key ideas in the theory of contextual integrity, its differences from existing approaches, and its harmony with basic intuition about information sharing practices and norms.

Keywords Privacy · Contextual integrity · Privacy law · Public policy · Protecting privacy · Networks

Introduction

In February 2012, the Obama White House unveiled a *Privacy Bill of Right*s (2012, 9), embedded in a comprehensive report, "Consumer Data Privacy in a Networked World: A Framework for Protecting Privacy and Promoting Innovation in the Global Digital Economy." In addition to the Bill of Rights, the Report's Framework

✉ Helen Nissenbaum
Helen.nissenbaum@nyu.edu

[1] Culture and Communication, NYU Media, 239 Greene St., 7th Fl., New York, NY 10003, USA

Published online: 12 July 2015

for Protecting Privacy laid out a multi-stakeholder process, articulated foundations for effective enforcement, pledged to draft new privacy legislation, and announced an intention to increase interoperability with international efforts (Civil 2012). The White House report was but one among several governmental studies and reports in the US and elsewhere (e.g. World Economic Forum (WEF) 2012; Federal Trade Commission (FTC) 2012) responding to increasingly vocal objections to information practices above and below the radar so out of control that in 2010 the *Wall Street Journal*, sentinel of business and commercial interests, launched a landmark investigative series *What They Know*, which doggedly revealed to readers remarkable and chilling activities ranging from ubiquitous online monitoring to license plate tracking and much in between (Angwin and Valentino-Devries 2012; Valentino-Devries and Singer-Vine 2012). The dockets of public interest advocacy organizations were filled with privacy challenges. Courts and regulatory bodies were awash with cases of overreaching standard practices, embarrassing gaffes, and technical loopholes that enabled surreptitious surveillance and the capture, aggregation, use, and dispersion of personal information.

As awareness spread so did annoyance, outrage, and alarm among ordinary, unsophisticated users of digital and information technologies as they learned of practices such as Web tracking, behavioral advertising, surveillance of mobile communications, information capture by mobile apps (including location), capture of latent and revealed social network activity, and big data.[1] [It bears mentioning that although rhetoric often names the technologies themselves as sources of concern, e.g. "big data," or "biometrics," the sources of privacy threats are socio-technical systems, that is to say, technologies embedded in particular environments shaped by social, economic, and political factors and practices and put to specific purposes (Nissenbaum 2010).] Most salient to individuals are practices of familiar actors with which they are directly acquainted, such as Facebook, Google, Amazon, Yelp, and Apple. More informed critics point to information brokers, backend information services, ad networks, voter profilers, "smart grids," surveillance cameras, and biometric identification systems, to name a few, which relentlessly monitor and shape lives in ways neither perceptible nor remotely comprehensible to the public of ordinary citizens.

Acknowledging the problem, governmental bodies in the US have kept citizens' privacy on the active agenda, pursuing cases against specific activities (e.g. Google Inc. v. Joffee et al. 2014; FTC v. Wyndham Worldwide Corporation, et al. 2014; re: Netflix Privacy Litigation 2012). They have conducted studies, public hearings, and multistakeholder deliberations on specific practices, such as commercial uses of facial recognition systems, surreptitious uses of personal information by mobile apps, and applications of big data (National Telecommunications and Information Administration (NTIA) 2013a). Such initiatives are also underway in Europe in governmental as well as nongovernmental sectors, including, for example, the World Economic Forum, the Organization for Economic Co-operation and

[1] Anxiety over the digital age, and more specifically, big data, is a major theme in mainstream tech and business journalism as of 2013. For more information, see *The New York Times'* special section "Big Data 2013." http://bits.blogs.nytimes.com/category/big-data-2013/.

Development (OECD), and the European Commission (WEF 2012; European Union 2013).

For those who have followed academic and public deliberations, these cycles of revelation and protest are not new. A more pointed source of incredulity, however, is that this panoply of information practices, for the most part, proceeds under the halo of legality, quite literally, evoking gasps of disbelief among the newly informed. For privacy scholars and activists, the level of indignation about these perfectly lawful practices adds strength to their position that something is amiss in the relevant bodies of law and regulation—the status quo needs correction. In the recent history of privacy, the present moment resembles others, in which new technologies, practices, or institutions are seen to cross a threshold, setting off a cry that, "Something has to be done!" (Ellul and Merton 1964; Ware 1967; Brooks 1980; Regan 1995).

Responding to this call, the White House and FTC reports outlined ways of reaching beyond the status quo. This paper focuses on the White House Consumer Privacy Bill of Rights and within it, the Principle of Respect for Context (PRFC). It argues that how this Principle is interpreted is critical to the success of the Privacy Bill of Rights as an engine of change—whether it succeeds in its mission of change or devolves to business as usual.

White House Report and Respect for Context

Until the Department of Commerce took up its study of privacy, a prelude to the 2012 White House Report, the FTC had been the key government agency spearheading important privacy initiatives in the commercial arena with rulemaking and legal action. The report signaled direct White House interest in contemporary privacy problems and buoyed hopes that change might be in the air. The Report and Bill of Rights were cautiously endorsed by a range of parties who have disagreed with one another on virtually everything else to do with privacy. On the public interest advocacy front, the Electronic Frontier Foundation, for example, which had proposed its own Bill of Privacy Rights for Social Network Users, conceded that, "this user-centered approach to privacy protection is a solid one" (Hoffman 2012). The Electronic Privacy Information Center (EPIC) "praised the framework and the President's support for privacy, and said that the challenge ahead would be implementation and enforcement" (EPIC.org 2012), and The Center for Democracy and Technology (CDT) "welcome[d] the Administration's unveiling," endorsing the report's "call for the development of consensus rules on emerging privacy issues to be worked out by industry, civil society, and regulators" (CDT 2012). On the industry front, Google declared itself, "on board with Obama's Privacy Bill of Rights," and Intel affirmed the Administration's " … calls for US federal privacy legislation based upon the Fair Information Practices," (Hoffman 2012). Chris Civil, in an overview of the bill and reactions to it, cited *Time's* observation that, "the most "remarkable" element of the new framework is that it has not been greeted with outrage from Silicon Valley companies, who have previously opposed similar privacy legislation efforts led by the California State Senate" (Civil 2012).

Of the seven principles proposed in the Consumer Privacy Bill of Rights, six are recognizable as kin of traditional fair information practice principles, embodied, for example, in the *OECD Privacy Guidelines*.[2] However, the PRFC, the expectation that "companies will collect, use, and disclose personal data in ways that are consistent with the context in which consumers provide the data," (p. 47) is intriguingly novel and, in part, a reason why the Report suggested that something beyond business-as-usual was its aim. How far the general endorsement of respect-for-context will push genuine progress, however, is critically dependent on how this principle is interpreted. Context is a mercilessly ambiguous term with potential to be all things to all people. Its meanings range from the colloquial and general to the theorized and specific, from the banal to the exotic, the abstract to the concrete, and shades in between. If determining the meaning of context were not challenging enough, determining what it means to respect it opens further avenues of ambiguity. In short, the positive convergence of views held by longstanding antagonists may be too good to be true if it rests on divergent interpretations. Whether the Privacy Bill of Rights fulfills its promise as a watershed for privacy and whether the principle of respect for context is an active ingredient in the momentum will depend on which one of these interpretations drives public or private regulators to action.

Meanings of Context

Setting aside general and colloquial uses, as well as idiosyncratic ones, this article takes its cues from specific meanings and shades of meanings embodied in recorded deliberations leading up to the public release of the Report and in action and commentary that has followed it, all clearly influential in shaping the principle. My purpose, however, is not purely semantic; it does not involve judging some of these meanings to be "correct" while others "incorrect." Instead, it is to highlight how different meanings imply different policy avenues, some seeming to favor the entrenched status quo, others to support progressive if limited improvement. Ultimately, I will argue that the interpretation that opens doors to a genuine advancement in the policy environment is embodied in the theory of contextual integrity: it heeds the call for innovation, recognizes business interests of commercial actors, and at the same time places appropriate constraints on personal information flows for the sake of privacy. The paper does not argue that it is incorrect to use context in the myriad of ways we do, merely that only a subset of uses systematically favor certain policy directions over others, and, more importantly, not all among this subset promise a productive departure from "business as usual."

In the subset of interpretations with systematic implications for policy, four are of particular interest because they reflect persistent voices in discussions leading up to and following the White House report: context as technology platform or system, context as sector or industry, context as business model or practice, and context as

[2] The remaining six principles are Individual Control, Transparency, Security, Access and Accuracy, Focused Collection and Accountability.

social domain. Although within each of the four there are nuances of meaning and subtleties of usage, for purposes of this discussion, they have been set aside or, where possible, absorbed into the core. One example of this is *the context of a relationship*, which is more general and abstract than the four listed. In deciding whether this framing warranted a separate analysis, I examined comments from the Online Publishers Association introducing this phrase. Finding that it was referring specifically to the relationship between publishers and their clients (readers, viewers, etc.), I was comfortable absorbing this understanding of context within that of business practice. After considering all four interpretations, this paper contends, ultimately, that social domain is the only interpretation of context that marks a meaningful departure from business as usual.

There are many ways context may be relevant to those modeling human behavior. In explaining online behavior, for example, contextual factors such as geo-location, time, stage in a series, or a myriad other possibilities may be external to a given model but may serve to refine its descriptive or predictive accuracy, helping to explain and predict at finer grain behaviors such as web search, receptiveness to advertising, and even vulnerability to malevolent overtures, such as phishing attacks (Kiseleva et al. 2013a, b). Understood in this way, as a factor external to a given model, this general sense of context can also serve to refine empirical theories of privacy. One can observe that expectations are affected by the context of a promise, a relationship, a conversation, or an event. Place—geo-spatial or physical location—such as, home, office, café, supermarket, park, corner of Broadway and Bleecker, is a particularly salient contextual refinement (see, e.g. Dwork and Mulligan 2013). Context as place is of natural interest not only because it reflects common English usage, but also because, historically, it has served to qualify privacy expectations, such as in distinguishing the home from public space (US Const. amend. IV; Selbst 2013).

I have not given independent consideration to context abstractly conceived because I have not seen systematic ties to specific expectations of privacy. Although place is a significant factor in accounting for privacy expectations, it was not singled out in the White House Report. Although, undoubtedly, place is important in shaping privacy expectations it does so not necessarily *as an independent factor*, that is, whether an activity takes place inside a building or outside, at one particular geo-location or another, but as it functions in social terms, as, say, a church, home, or hospital—as will be clarified later in this article.

Context as Technology System or Platform

Many privacy issues we are confronting emerge from the realm of digital networks—the Internet, and the myriad platforms and systems sitting atop (or below) it, such as, mobile systems, email, social networks, cloud providers, and the Web itself. For most of us these disparate technical substrates, systems, and platforms are experienced indistinguishably from one another and, though technical experts give a more rigorous account of their differences, they are akin, from the perspective of user experience and political economy. We talk of communication and transaction taking place *online* or *in* Cyberspace, and the privacy problems

emerging from them are associated with these electronically mediated contexts without a clear sense that they may emerge in different ways because of the different architectures and protocols. They become the problems of online privacy—problems of a distinctive domain requiring a distinctive approach. It is a short distance to conceive of this technological substrate as a context, one that makes a difference to privacy; we readily conceive of talking in the context of, say, a phone call, acting in the context of an online social network, expressing ourselves in the contexts of Twitter, Facebook, and Wikipedia, or in the contexts of a mobile app, or location-based services. In such expressions contexts are defined by the properties of respective media, systems, or platforms whose distinctive material characteristics shape—moderate, magnify, enable—the character of the activities, transactions, and interactions they mediate. They also shape the ways information about us is tracked, gathered, analyzed, and disseminated. If properties of technical systems and platforms define contexts, then a principle that supports *respect* for contexts presumably implies that policies should be heedful of these defining properties of systems and platforms.

The idea of context as technical system or platform is suggested in the Foreword of the White House report when it states:

"Privacy protections are critical to maintaining consumer trust in networked technologies. When consumers provide information about themselves—whether it is in the context of an online social network that is open to public view or a transaction involving sensitive personal data—they reasonably expect companies to use this information in ways that are consistent with the surrounding context. Many companies live up to these expectations, but some do not. Neither consumers nor companies have a clear set of ground rules to apply in the commercial arena. As a result, it is difficult today for consumers to assess whether a company's privacy practices warrant their trust" (White House Privacy Report 2012, 5).

Comments by others reflect a similar interpretation. AT&T, for example, notes that diverse technical platforms generate distinctive challenges to privacy: "Indeed, the power of Web 2.0 inter-related media is precisely that content can be used in ways that were not expected or understood when they were collected" (Intel 2011, 8). Google encourages enforceable codes of conduct that, "reflect changing practices, technologies and shifting consumer expectations" (Intel 2011, 9); and Intuit observes that, "Collecting information for use in routing a request on the Internet should have different standards for transparency, acceptable uses, protection, and retention than the information collected to describe a patient's visit to a physician …" (Intel 2011, 11). Finally, the idea that technology defines context is suggested in the framing of NTIA's July 2012 kickoff multi-stakeholder (MSH) process around mobile applications, suggesting that mobile apps define a normative category.[3]

[3] July, 2012. *Multistakeholder Process To Develop Consumer Data Privacy Code of Conduct Concerning Mobile Application Transparency*. Symposium conducted at the open meeting of The National Telecommunications and Information Administration, Washington, DC.

Context as Business Model or Business Practice

In the discourse surrounding the Report the interpretation of context as prevailing business model or business practice was evident in various comments, particularly those offered by incumbents in the IT and information industries, for example, "Technology neutral and flexible legislation can actually aid small business growth as it provides a clear set of 'rules of the road' for everyone, while at the same time allowing those rules to be adapted to each business' unique situation," (Intel 2011, 4). This comment suggests that technology per se does not define privacy rules of the road but that these should be guided by the needs of distinctive business models aimed at promoting growth. Similarly, "TRUSTe supports the continued role of industry in defining purpose specifications and use limitations based on the unique needs of a company's business model" (Maier 2010, 8). According to Google, "The fast-paced introduction of new Internet services drives equally rapid shifts in consumer expectations and preferences. An effective privacy regime must allow for real time reactions to address changes in consumer privacy preferences resulting from the introduction and adoption of new tools and services" (Chavez 2011, 2). AT&T urges, "This flexibility should also allow companies to describe the use of data within broad categories, such as 'for marketing purposes,' without the need specify the particular purpose for the collection of each piece of data. Indeed, the power of Web 2.0 inter-related media is precisely that content can be used in ways that were not expected or understood when they were collected" (Raul et al. 2011, 17). Asserting a special privilege for the business practices of online publishers, the Online Publishers Association (OPA), with members including *WebMD*, *FoxNews*, and *The New York Times*, claims that, "Online publishers share a direct and trusted relationship with visitors to their websites. In the context of this relationship, OPA members sometimes collect and use information to target and deliver the online advertising that subsidizes production of quality digital content" (Horan 2011, 6).

Interpreted as the model or practice of a particular business, context is established according to that business' aims and the means it chooses to achieve these aims. There is nothing surprising about merchants orienting their buying and selling practices around profitability, so we should not be surprised that information service providers orient their models around growth and competitive edge. According to this understanding, contexts are defined by particular business models, in turn shaping respective information flow practices. Taking Google's comment above as a concrete case-in-point, this interpretation suggests that contexts generated by its business-driven Internet services, for example, shape consumer expectations of privacy, and not the other way around. Similarly, AT&T speculates that the privacy assumptions users hold will bend flexibly to the contours of "marketing purposes," defined as whatever is needed to strengthen a business model.

Context as Sector or Industry

Endorsing the sectoral approach that the United States has taken to privacy protection, TRUSTe notes that, "the regulatory frameworks currently in place in the US reflect this inherently contextual nature of privacy e.g. Fair Credit Reporting Act

(FCRA)/Fair and Accurate Credit Transactions Act (FACTA) (information used in "consumer reports"), Gramm-Leach-Bliley (information sharing between financial institutions and affiliates), Health Insurance Portability and Accountability Act (HIPAA) (transactions involving protected health information by "covered entities")" (Maier 2010, 2). In a similar vein: "Intuit's experience in multiple sectors has taught us that providers and consumers of information in the health sector, for example, have different requirements and expectations for protection than do those in financial services. … Subject matter experts could help inform the development of appropriately balanced codes" (Lawler 2011, 9).

I have placed "industry" in the same category as "sector," not because they have identical meanings, but because, in practice, these terms are used interchangeably in the commentaries from which I rendered the category. Adopting the interpretation of context as sector or industry, respect for context would amount to adherence to the set of rules or norms developed by, for and within respective sectors or industries.

Context as Social Domain

This interpretation, supported by the theory of contextual integrity, presents contexts as social spheres, as constituents of a differentiated social space. As such, they serve as organizing principles for expectations of privacy. Although contextual integrity relies on an intuitive notion of social sphere, covering such instances as education, healthcare, politics, commerce, religion, family and home life, recreation, marketplace, work and more, scholarly works in social theory and philosophy have rigorously developed the concept of differentiated social space, though with diverse theoretical underpinnings and terminology (e.g. sphere, domain, institution, field[4]). In intuitive as well as academic accounts, spheres generally comprise a number of constituents, such as characteristic activities and practices, functions (or roles), aims, purposes, institutional structure, values, and action-governing norms. Contextual norms may be explicitly expressed in rules or laws or implicitly embodied in convention, practice, or merely conceptions of "normal" behavior. A common thesis in most accounts is that spheres are characterized by distinctive internal structures, ontologies, teleologies, and norms.

From the landscape of differentiated social spheres, the theory of privacy as contextual integrity develops a definition of informational privacy as well as an account of its importance. Taking context to mean social sphere, respect for context would mean respect for social sphere. To explain what *this* means and why it opens new and significant avenues for the proposed White House policy framework requires a brief excursus into the theory of contextual integrity.

A Detour: Theory of Contextual Integrity

Other accounts of the profound anxiety over privacy, fuelled by the steep rise in capture, analysis, and dissemination of personal information, point to the loss of

[4] For a further discussion on spheres, see Nissenbaum (2010 pp. 80, 131, 166–169, 198–200, 240–241).

control by data subjects and sheer increased exposure. Although these factors are part of the story, the theory of contextual integrity holds the source of this anxiety to be neither in control nor secrecy, but appropriateness. Specifically, technologies, systems, and practices that disturb our sense of privacy are those that have resulted in *inappropriate* flows of personal information. Inappropriate information flows are those that violate context specific informational norms (from hereon, "informational norms"), a subclass of general norms governing respective social contexts.

Aiming at descriptive accuracy, the theory articulates a model wherein informational norms are defined by three key parameters: information types, actors, and transmission principles. It postulates that whether a particular flow, or transmission of information from one party to another is appropriate depends on these three parameters, namely, the type of information in question, about whom it is, by whom and to whom it is transmitted, and conditions or constraints under which this transmission takes place. Asserting that informational norms are context-relative, or context-specific, means that within the model of a differentiated social world, they cluster around and function according to coherent but distinct social contexts. The parameters, too, range over distinct clusters of variables defined, to a large extent, by respective social contexts.

Actors—subject, sender, recipient—range over context relevant functions, or roles, that is, actors functioning in certain capacities associated with certain contexts. These capacities (or functional roles) include the familiar—physician, nurse, patient, teacher, senator, voter, polling station volunteer, mother, friend, uncle, priest, merchant, customer, congregant, policeman, judge, and, of course, many more. In complex, hierarchical societies, such as the contemporary United States, actors governed by informational norms might be collectives, including institutions, corporations, or clubs.

The parameter of information type, likewise, ranges over variables derived from the ontologies of specific domains. In healthcare, these could include symptomologies, medical diagnoses, diseases, pharmacological drugs; in education, they may include cognitive aptitude, performance measures, learning outcomes; in politics, party affiliations, votes cast, donations; and so forth. There are, in addition, types of information that range across many contexts, to give a few basic examples, name, address, and gender.

Transmission principle, the third parameter, designates the terms, or constraints under which information flows. Think of it as a sluicegate. Imagine that you are applying for a bank mortgage on a new home and have signed a waiver allowing the bank to obtain a copy of your credit report from Equifax. To map this transaction onto the structure of context specific informational norms: (1) actors: you, the applicant, are the data subject; the bank is the data recipient; and the credit bureau is the sender; (2) information type includes the various fields of information that are provided in a credit report; and (3) transmission principle, is "with the information subject's signed waiver." The transmission principle, abstractly conceived, has not been explicitly recognized in scholarly or policy deliberations even though, in practice, its implicit role in social convention, regulation and law can be pivotal. Isolating the transmissions principle as an independent variable also offers a more general account of the dominant view of a right to privacy as a right to control

information about ourselves. Through the lens of contextual integrity, this view mistakes one aspect of the right for its entirety, for control over information by the information subject is but one among an extensive range of possible transmission principles, including, "in confidence," "with third-party authorization," "as required by law," "bought," "sold," "reciprocal," and "authenticated," among others.

A feature of informational norms that bears emphasizing is that the three parameters—actors, information types, and transmission principles—are independent. None can be reduced to the other two, nor can any one of them carry the full burden of defining privacy expectations.[5] This is why past efforts to reduce privacy to a particular class of information—say "sensitive" information—or to one transmission principle—say, control over information—are doomed to fail and, in my view, for decades have invited ambiguity and confusion, hindering progress in our understanding of privacy and attempts to regulate its protection. Control over information is an important transmission principle, but always with respect to particular actors and particular information types, all specified against the backdrop of a particular social context. Although much could be said about each of the parameters, the scope of this paper limits us.[6]

Contextual integrity is achieved when actions and practices comport with informational norms. But when actions or practices defy expectations by disrupting entrenched, or normative information flows, they violate contextual integrity. As such, informational norms model privacy expectations. When we find people reacting with surprise, annoyance, and indignation, protesting that their privacy has been compromised, the theory would suggest as a likely explanation that informational norms had been contravened, that contextual integrity had been violated. Conversely, informational norms may serve as a diagnostic tool with prima facie explanatory and predictive capacities. From observations of technical systems or practices, which result in novel patterns of information flow according to actors, information types, or transmission principles, the theory would predict that people may react with surprise and possibly annoyance. Contextual integrity provides a more highly calibrated view of factors relevant to privacy than traditional dichotomies such as disclose/not disclose, private/public.

The diagnostic or descriptive role of contextual integrity is not the full story, but before turning to the ethical dimension, two quick implications bear mentioning. One is that when it comes to the nuts and bolts of privacy law, policy, and design, area experts in respective contexts—education, healthcare, and family and home-life—are crucial to understanding roles, functions, and information types. They, not privacy experts, are best equipped to inform processes of norm discovery, articulation and formation. A second implication is that though practices in well-circumscribed social institutions may be thickly covered by informational rules, only a fraction of all possible information flows in daily life are likely to be covered by explicit norms. Compare, for example, a court of law, a stock exchange and a

[5] In practice, we may omit explicit mention of one or two of the parameters where these are obviously understood, or tedious to fully specify.

[6] Greater detail can be found in *Privacy in Context* (Nissenbaum 2010), however, the role and scope of transmission principles deserves even fuller coverage elsewhere.

hospital with an informal social gathering, a shopping mall, a beauty parlor—picking a few at random. The lens of contextual integrity provides a view of emerging digital (sociotechnical) information systems in terms of radical disruptive information flows, in turn an explanation of contemporary anxiety and acute concern over privacy. But many novel information flows are disruptive not because they contravene explicit norms, but because they open up previously impossible (possibly unimaginable) flows. In these instances, consternation follows because flows are unprecedented, may or may not expose new vulnerabilities and hazards. How to cope with these puzzling cases, in addition to the ones in which existing norms are violated, are challenges for the prescriptive dimension of contextual integrity.

Contextual Integrity: Ethics and Policy

Novelty and disruption are not problematic even if they result in direct contraventions of entrenched informational norms. Even a superficial survey reveals many welcome alterations in flows brought about by adoption of information and network technologies; for example, enhanced health indicators, robust and cheap new forms of communication and association, such as through social networks, and information search tools online. In many of these instances novel flows have replaced suboptimal ones that had become entrenched in particular contexts due to the limits of past technologies, media, or social systems.[7] Questions must be addressed, however. How to evaluate disruptive information flows brought about by novel technologies, media, and social systems; how to distinguish those that embody positive opportunities from those that do not; those that violate privacy from those that do not—all important challenges for any theory of privacy. When A.T.&T. asserts, "Consumers approach the Internet with a consistent set of expectations, and they should be able to traverse the Internet having those expectations respected and enforced" (2012), it endorses the normative clout of our privacy expectations. And because we may not agree that *all* expectations deserve to be met, we can reasonably require of a theory of privacy to account for the difference between those that do and those that do not. This is the challenge any normative theory of privacy should address and it is the challenge for which a normative dimension of contextual integrity was developed.

A fundamental insight of contextual integrity is that because information flows may systematically affect interests and realization of societal values, these can be used as touchstones for normative evaluation. Where novel flows challenge entrenched informational norms, the model calls for a comparative assessment of entrenched flows against novel ones. An assessment in terms of interests and values involves three layers. In the first, it requires a study of how novel flows affect the interests of key affected parties: the benefits they enjoy, the costs and risks they suffer. These may include material costs and benefits as well as those less palpable, including shifts in relative power. Beyond this largely economic analysis, frequently followed in policy circles, the normative analysis directs us to consider general

[7] For development of this point, see Nissenbaum (2010).

moral, social, and political values. These would include not only costs and benefits but also considerations of fairness, the distribution of these costs and benefits, who enjoys the benefits and who endures the costs. Thus, for example, where new flows involve power shifts, this second layer asks whether the shifts are fair and just. Other core ethical and societal values that have been identified in a deep and extensive privacy literature are democracy, unfair discrimination, informational harm, equal treatment, reputation, and civil liberties. This literature has shone light particularly on the connections between privacy and aspects of individual autonomy, including moral autonomy, boundary management, and identity formation.[8]

The third layer introduces a further set of considerations, namely, context-specific values, ends, and purposes. This layer sets contextual integrity apart from many other privacy theories.[9] It offers a systematic approach to resolving conflicts among alternative patterns of information flows, which serve competing interests and values respectively. In a particular context, one pattern of flow might support individual freedom, an alternative, safety and security. The additional analytic layer may resolve the conflict. In some, freedom will trump, in others, security will trump depending on facts on the ground and respective goals and values. Although privacy is often pitted against the interests of business incumbents, or is viewed as conflicting with values such as national security, public safety, and freedom of expression, contextual integrity allows us to unravel and challenge such claims. This layer insists that privacy, as appropriate information flows, serves not merely the interests of individual information subjects, but also contextual, social ends and values.

In the context of healthcare, for example, where the integration of electronic patient records has radically altered flows of information, it is crucial to ask how these have affected the attainment of ends and purposes of healthcare and whether the values associated with healthcare are sustained. In the US, these aims might include curing and preventing disease, repairing bodily injury, and minimizing physical pain, while values include patient autonomy, frugality, equal access, and non-discrimination. Thus, when assessing terms of access to medical records, although patient interests (including freedom from embarrassment and shame) are an important consideration, as are the interests of other stakeholders, an analysis must also consider the purposes and values of the healthcare context. If individuals avoid diagnosis and treatment because of access rules that are too lax, not only do they suffer but others, too, pay the price, and ends and values of healthcare are undermined. A similar argument explains why ballot secrecy, or privacy, is crucial in democratic elections: it not only protects individual voters against intimidation and, possibly retribution, but promotes democracy itself, which is based on autonomous preferences of individual citizens.[10]

[8] For example, see Cohen (2012), Solove (2006), Van den Hoven (1998), Schoeman (1984) and Gavison (1980).

[9] Nissenbaum (2010).

[10] For a longer, more elaborate discussion, see Nissenbaum (2010, 2011, 2012).

The claim of this paper is that context, understood as social sphere, is far more likely to yield positive momentum and meaningful progress in privacy law and policy than understood as technology, sector, or business model. With context-specific informational norms establishing the link between context and privacy, *respect for context* amounts to respect for contextual integrity. To flesh out this claim, a fresh look at the White House Privacy Bill of Rights will be instructive.

Respect for Context and the Consumer Internet Privacy Bill of Rights

The White House Privacy Bill of Rights embodies "fair information practice principles" (FIPPS), as have many codes of privacy before it, in the US and internationally. Appendix B of the report accounts for its debt to FIPPS and other codes in a table that lines up respective principles of the Consumer Privacy Bill of Rights (CPBR) alongside respective principles in the OECD Privacy Guidelines, the Department of Homeland Security (DHS) Privacy Policy, and Asia–Pacific Economic Cooperation (APEC) Principles (2012, 59).[11] The CPBR principles of Transparency, Security, Access and Accuracy, and Accountability have relatively straightforward counterparts in the other sets of guidelines each worthy, in its own right, of in-depth critical analysis. Respect for Context, the focus of this article, is aligned with Purpose Specification and Use Limitation Principles. The White House's CPBR principles of Focused Collection and Individual Control, whose counterparts in the OECD Guidelines are listed as Collection and Use Limitation principles, would therefore also be affected by the interpretation of Context.

Let us zoom in for a closer look at the right of Respect for Context, "a right to expect that companies will collect, use, and disclose personal data in ways that are consistent with the context in which consumers provide the data" (White House Privacy Report 2012, 55). Its close kin, given as, (1) Purpose Specification and (2) Use Limitation, require that, (1) "The purposes for which personal data are collected should be specified no later than at the time of data collection and the subsequent use limited to the fulfillment of those purposes or such others as are not incompatible with these purposes and as are specified on each occasion of change of purpose" (p. 58); and (2) "Personal data should not be disclosed, made available or otherwise used for purposes other than those specified in accordance with Paragraph 9 (i.e. purpose specification) except … (a) with the consent of the data subject; or (b) by the authority of law" (p. 58).

Speaking philosophically, we can say that the Purpose Specification and Use Limitation principles have only indexical meaning, emerging in particular, concrete instances of use. Once purposes are specified uses, too, are limited accordingly. But what these purposes are, or may be, is not given in the principles themselves. One could admire the adaptability of these principles—a virtue of FIPPs, by some counts. Or, point out, as Fred Cate has, that FIPPS themselves do not provide privacy protection, merely procedural guidance whose substantive clout is

[11] "Appendix B: Comparison of the Consumer Privacy Bill of Rights to Other Statements of the Fair Information Practice Principles (FIPPS)," White House Privacy Report 2012.

indeterminate.[12] According to Cate the FIPPS purpose specification principle offers some traction for privacy protection. He points out, however, that unless constraints are placed on what purposes are legitimate (and why), a purely procedural Purpose Specification Principle opens a glaring loophole in FIPPS.[13] This point is crucial for my argument about context.

Use Limitation, in turn, is compromised by the wild-card character of Purpose Specification, as is the principle of Collection Limitation (often called Data Minimization), which restricts information collection to that which is necessary for specified purposes. Talk about a vicious circle! Other principles which may seem to be inoculated against this indexicality are also affected, albeit indirectly. Take Security and Data Quality requirements. Although no explicit mention is made of purpose in these principles they are implied, as what counts as reasonable standards for both is surely a function of the purposes for which information is gathered and for which it is earmarked—e.g. whether the information in question is being collected for purposes of national security versus consumer marketing. The meaning of these principles is dependent on purpose, and the data collector, at will, may specify purpose. Unless and until purposes are shaped by substantive requirements, FIPPS constitutes a mere shell, formally defining relationships among the principles and laying out procedural steps to guide information flows. Given the centrality of FIPPs in virtually all privacy (or data protection) policies throughout the world, it is surprising to find that privacy is elusive, and even that fairness itself can be questioned in the contemporary regimes of privacy policies (Nissenbaum 2011).

A Question of Interpretation

The rhetoric surrounding NTIA's release of the Consumer Privacy Bill of Rights (CPBR) was that of turning a new page, ambitious and optimistic. The principle of Respect for Context offered a salient departure from FIPPs' Purpose Specification and Use Limitation principles. Herein lay the promise of something materially different, something better. But whether the promise can be fulfilled and not devolve to business as usual will depend on how we interpret context. In the previous section, we saw that the interpretation of Respect for Context is important not only in its own right, but is pivotal, too, for fixing meanings for other key principles, including Access and Accuracy, Focused Collection, and Security. Fixing meanings *correctly*, that is, in a way that the innovation embodied in Respect for Context materially advances the state of privacy protection in the US is, therefore, critical. Below, I will explain why, among the four alternatives, context understood as social domain is the most viable basis for progress.

Consider context as business model or practice. Under this interpretation, context would be determined by the exigencies of a particular business and communicated

[12] For Cate's cogent analysis, see Cate (2006). See another astute discussion in Rubinstein (2010).

[13] In fairness, others in the policy arena have noted the indeterminacy of the linchpin purpose specification and use limitation principles and are attempting to set substantive standards. For example, the EU Article 29 Working Party in Opinion 03/201d on purpose limitation and aspects of the problem discussed in Rauhofer (2013).

to individuals via general terms of service. In the context of an online purchase of physical goods, for example, it is reasonable for a merchant to require a consumer's address and valid payment information. But if the business purpose is a blank check, we are in trouble. Even in this simple illustration, questions remain: What happens to the information after delivery is completed? With whom can this information be shared, and under what terms? For how long, and who is responsible if harm follows its unintended leakage, or theft by criminals? With the ever-growing thirst for data, questions such as these have multiplied by orders of magnitude and while our intuitions are robust when it comes to merchants of physical goods, reasonable purpose for businesses *in* the information business is murkier still.

If business model and practice define context, political economy would shape the relationship between the information collector and information subject allowing no recourse to standards beyond business expedience (except in the few sectors where privacy legislation exists). By definition, each business entity determines what is and is not expedient. Other standards, such as security, use limitation, collection minimization, and access, which all are defined in terms of purpose, will be defined accordingly. Defining context as business model leaves the door wide open to anything reasonably conceived as profitable for respective businesses—buying up information resources, extracting information resources from transactions, and using such information in any manner (limited only by positive law and regulation.) This is not to say that business models are irrelevant to context and informational norms, only that the promise of change will not be fulfilled if business interests are the sole arbiters of context (Friedman 1970). Although business needs are an important consideration, they do not form a sound basis for privacy's moral imperative.

What about context as technology platform or system? First, consider what this means. It is quite sensible to refer to a Facebook profile, a Bing search, a Fitbit group, the Web, an email exchange, and a Google + Hangout as contexts. The question here, however, is not whether it is *sensible* to use the term context in these ways but whether these ways can form the reference point for Respect for Context. Answering affirmatively means technological affordance would determine moral imperative; it means accepting that whatever information flows happen to be afforded by a social network, a Web search engine, health-tracking device, and so forth, not only determine what *can* happen but what *ought* to happen. In these stark terms, the thesis may seem absurdly counterintuitive, yet it is embodied in familiar practices and reasoning. Take, for example, controversies surrounding online tracking. After conceding there was strong support for providing to individuals the means to delete third party cookies, various workarounds emerged, such as flash cookies and browser fingerprinting that reinstated cross-site tracking functionality. If technological affordance defines moral imperative there are no grounds for critiquing the workarounds. Similarly, when Mark Zuckerberg stated that Facebook had altered norms because the system had altered actual flows, he was right, by definition, because whatever flows are enabled by platforms simply *are* the flows that context legitimates.

Denying that technological affordance defines respect for context does not mean it is irrelevant to it. Practices are changed and sometimes they pull norms and standards along with them. The explosive growth of socio-technical information systems, the source of much consternation over privacy, is responsible for radical

disruptions in information gathering, analysis, and distribution, in the types of information that are accessed, analyzed and distributed, the actors sending and receiving information, and in the constraints or conditions under which it flows. These disruptions not only divert information flows from one path to another and one recipient to another, or others, but also may reconfigure ontologies, yield new categories of information, and new types of actors and modes of dissemination. Such changes may call for the reconsideration of entrenched norms and development of norms where none previously may have existed.

The "old" technologies of the telephone, for example, introduced novel parameters of voice dissemination including new classes of actors, such as telecommunications companies, in the early days, human operators, later on, mechanical and electronic switches. Existing norms of flow governing communications and, say, eavesdropping, may provide initial models for new conditions afforded by the telephone. As novel systems cause increasing divergence from pre-existing affordances, novel challenges demand deeper examination of what is at stake in a social world whose transactions, conversations, and relationships have been reconfigured by telephonic media. A pair of famous US Supreme Court cases, roughly 40 years apart, reveal this progression: Olmstead v. United States, 277 US 438 (1928) and Katz v. United States, 389 US 347 (1967). Landmark Fourth Amendment cases involving an historical reversal of law, these cases have been endlessly analyzed and taught. The common lesson drawn from them, which I have no cause to challenge, is that the 1967 Court finally "got it right." Shifting attention from the foreground of what counts as a legitimate expectation of privacy, to the background, of how the world had changed, we note that as telephones became normalized, phone-mediated conversations became integral to social life. In my view, this is key to explaining *why* the Court "got it right" in the Katz case. The ascent of telecommunication in social, political and economic life also meant addressing head-on the status of newly emerging actors, forms of information, and constraints on flow. To this day (underscored by the Snowden revelations) we are living with the consequences of legislation that attempted to define duties of phone companies, and the varied types and degrees of access they (and others) would have to the new forms of data generated by the telephonic medium, from pen register data to content of phone calls.[14]

Technical systems and platforms shape human activity by constraining and affording what we can do and say; in this sense, they are rightly conceived as contexts and deserve to be objects of attention and regulation. Allowing that people act and transact in contexts shaped by technical systems does not mean, however, that these systems fully account for the meaning of Respect for Context. So doing

[14] 18 USC § 2511(2)(a)(i) 2011, (i): "It shall not be unlawful under this chapter for an operator of a switchboard, or an officer, employee, or agent of a provider of wire or electronic communication service, whose facilities are used in the transmission of a wire or electronic communication, to intercept, disclose, or use that communication in the normal course of his employment while engaged in any activity which is a necessary incident to the rendition of his service or to the protection of the rights or property of the provider of that service, except that a provider of wire communication service to the public shall not utilize service observing or random monitoring except for mechanical or service quality control checks." Thanks to Chris Hoofnagle for calling attention to this crucial point.

allows material design to define ethical and political precepts; it allows the powers that shape the technical platforms of our mediated lives not only to affect our moral and political experiences through built-in constraints and affordances, but further, to place them beyond the pale of normative judgment.

The practical implications of this distinction can be seen in relation to the first NTIA multistakeholder process. No fool's errand, its mission was to establish a code of conduct for mobile applications developers. The NTIA process, which (1) identified a new class of actors, including mobile app developers, among others and (2) articulated baseline constraints on appropriate behaviors in the ecologies of mobile information services, concluded with a set of guidelines (NTIA 2013b). In my view, respect for context, should not stop with these. Beyond the baseline, it would require that distinct sets of informational norms be fleshed out for mobile app developers according to the social meaning, or function, of their specific apps. Although developers of, say, Yelp, Google Maps, Four Square, Fitbit, and Uber should fulfill these baseline obligations in their collection, use, and disclosure of personal information, they do not stop with these. One could reasonably expect Fitbit to treat the information it gathers differently from, say Uber, or Four Square. Mobile app developers do not escape additional obligations of social context any more than physicians are relieved of duties of confidentiality when information is shared with them over the phone rather than during an office visit. Where technical platforms mediate multiple spheres of life the need to distinguish technological affordance from moral imperative is acute. Doubtless technologies shape contexts, may even constitute them, but where Respect for Context is a bellwether for privacy, it is a mistake to confuse technological contexts with those that define legitimate privacy expectations.

Interpreting context as sector or industry overcomes some of the drawbacks of context as business model, because instead of devolving to the self-serving policies of individual businesses, norms of information flow could be guided by a common mission of the collective—ideally, collective best practice. This interpretation also aligns with the US sectoral approach to privacy regulation and legislation, which, at its best, allows for the generation of rules that are sensitive to the distinctive contours of each sector. Extracting a Principle of Respect for Context, carrying moral weight, from a descriptive notion of sector requires a bridge. One is to recognize explicitly that sectors include more than industries, which range over a limited set of, primarily, business sectors. Existing practice in the US goes partway in this direction, in talk of education and healthcare, for example, as sectors. Extending the range to politics, family, or religion could deepen the appreciation of appropriate informational rules even further. Expanding and qualifying the scope of sectors in these ways, however, brings them close to the construct of social spheres around which the theory of contextual integrity is oriented.

Interpreting the Principle of Respect for Context as respect for contextual integrity means first, that any significant disruption in information flows triggers a call for analysis and evaluation in terms of types of information, actors, and transmission principles. Because shifts and changes characteristic of these disruptions may correspond to shifts and changes in the balance of power interests as well as achievement and abatement of values, identifying them is a crucial first step. Second,

an evaluation of disruptive flows extends beyond conventional measures of stakeholder interests and even beyond general moral and political values. It brings to the fore context-specific functions, purposes and values. Context is crucial to privacy, not only as a passive backdrop against which the interests of affected parties are measured, balanced, and traded off; rather, it contributes independent, substantive landmarks for *how* to take these interests and values into account. It makes the integrity of the contexts *themselves* the arbiter of privacy practices—vibrant marketplace, effective healthcare, sound education, truly democratic governance, and strong, trusting families and friendships.

Summary of Argument

For the Consumer Privacy Bill of Rights (CPBR) to advance privacy protection beyond its present state, a great deal hangs on how the Principle of Respect for Context (PRFC) is interpreted. Acknowledging the pivotal place context holds in the White House vision, commentaries have converged around four primary contenders: business model, technology, sector, and social domain. I have argued that respecting context as *business model* offers no prospect of advancement beyond the present state-of-affairs. Citing innovation and service as the drivers behind this interpretation, its proponents seem to expect individuals and regulators to sign a blank check to businesses, in collection, use, and disclosure of information based on exigencies of individual businesses.

Respecting context as *sector* (or industry) fares slightly better as it offers a framework beyond the needs of individual businesses for establishing standards and norms. How well this approach meaningfully advances privacy protection beyond the present state depends on how sectors are defined. If it follows the contours of industry, it might yield improvements in "best practices," but the interests of dominant incumbents may still prevail. This problem is particularly acute where the sector or industry in question is the "information sector," where the proverbial fox would be guarding the henhouse. Further, if industry dominates the construction of sectors, the influence of sectors such as healthcare, education, religion, politics, will be diminished, or the commercial aspects of these industries may play a disproportionate role. Correcting for these distortions brings sector-as-context closer to context as social domain. Understanding context in purely *technological* terms implies that legitimate expectations should be adjusted to reflect technical affordances and constraints. But so doing drains respect for context of moral legitimacy, getting things exactly backwards. Our morally legitimate expectations, shaped by context and other factors, should drive design and define the responsibilities of developers, not the other way around.

Interpreting context as *social domain,* as characterized in the theory of contextual integrity, avoids many of the problems associated with the other three. To respect context under this interpretation means to respect contextual integrity, and, in turn, to respect informational norms that promote general ethical and political values, as well as context specific ends, purposes, and values. Informational norms constitute the substantive cornerstone of policy and practice and replace both the serendipity

of design and arbitrary policies serving dominant parties. The ultimate contribution of contextual integrity does not rest with the concept of context, per se, but with two fundamental ideas behind it: One is the idea that privacy (or informational) norms require all relevant parameters to be specified including actors (functioning in roles), information types, and transmission principles. Omitting any one of these yields rules that are partial and ambiguous. The second fundamental idea is of context specific ends, purposes and values, which extend the significance of privacy beyond the balancing of interests, harms and benefits. Contextual integrity reveals the systematic dependencies of social values on appropriate information flows, once-and-for-all challenging the fallacy of privacy as valuable for individuals alone.

Conclusion: Implications for Practice

There are many meanings of the term context but not all of them are systematically tied to particular outcomes for privacy. In this paper, I have reviewed four that do make a difference and would lead to differences in the impact on privacy online of the Principle of Respect for Context in the White House Consumer Privacy Bill of Rights. To illustrate this claim on a concrete, if limited case, let us consider how the four interpretations might have played out in the case of 18 USC Section 2511 (2)(a)(i). As discussed above, 18 USC Section 2511 (2)(a)(i) prohibits telecommunications providers from intercepting, disclosing, or using the content of communications except in limited circumstances, such as, rendering service or protecting their property with further exceptions for legitimate needs of law enforcement and national security. How might the Principle of Respect for Context have shaped such legislation; specifically, what difference would the interpretation have made?

Let us begin with the interpretation of *context-as-technology*—not merely technology influencing context, but *defining* it. Under this interpretation, we would conclude that whatever interception, disclosure, or use of content is enabled by the mediating technologies should be "respected." Expectations of parties utilizing these technologies could not extend beyond what the technologies allow—for affordance defines legitimacy. Interpreted as *business model*, respect for context would allow individual providers to pursue whatever practices and policies they believe will promote profitability and an edge over competitors. These might include scanning conversations to pick out customers' commercially relevant interests, or providing access to interested parties willing to pay handsomely for access to conversations. I am not suggesting these particular outcomes are likely, merely the reasoning toward practice that this interpretation allows. Interpreting context as *sector* is likely to follow a slightly different track, if only because individual businesses, unless they collude, will seek to entrench practices that appeal to customers and level the playing field with competitors. Moreover, it is clear that how the boundaries, contours, and definition of sectoral groupings are set would affect policies and principles respective sectors support.

According to contextual integrity, interpreting context as social domain would focus attention on the role of telecommunications providers as communications' mediators. In this light, the tailored access rights devised by18 USC Section 2511

(2)(a)(i), allowing surveillance of conversations for the express purpose of assuring quality of service and protection of property, was a brilliant compromise. Laxer policies, as supported by the other interpretations, may discourage intimate or political conversation, as well as other sensitive conversations, such as, strategic business planning or path-breaking scientific collaborations, creating disadvantage for those needing to communicate securely and those benefitting from such communication. But beyond these impacts on various parties, they would reduce the utility of communications networks to individuals as well as their service of respective contextual ends, purposes, and values. Context as social domain draws attention to these higher order considerations, also reflected in the drafting of 18 USC Section 2511 (2)(a)(i).

As a brief aside, contextual thinking could have averted the Google Buzz fiasco.[15] Technological thinking may have suggested an alluring opportunity to leap frog into social networks based on Google's holdings from its email network. A business argument might have supported Buzz, in light of Facebook's success and Google's proprietary access to Gmail content and metadata. But failure to recognize that email serves multiple, disparate social contexts yielded an unappealing system and provoked outrage and indignation.

Contexts are shaped by technology, business practice, and industry sector. The may also be constituted by geographic location, relationship, place, space, agreement, culture, religion, era, and much more, besides. In individual cases, any of these factors could qualify and shape peoples' expectations of how information about us is gathered, used, and disseminated. No one of them, however, provides the right level of analysis, or carries the same moral and political weight as social domain. This is the thesis I have defended here. In light of it, I offer an amendment to the Consumer Privacy Bill of Right's Principle of Respect for Context:

> Respect for Context means consumers have a right to expect that companies will collect, use, and disclose personal data in ways that are consistent with the [social] context in which consumers provide the data.

Acknowledgments An early version of this paper was presented at the *Privacy Law Scholars Conference 2013* where James Rule, Mike Hintze, and other participants provided excellent commentary. I have benefitted from deep insights of many colleagues and from opportunities to present the work at the Amsterdam Privacy Conference, University of Washington, Fondation Télécom Seminar on The Futures of Privacy, and the EU JRC Ispra Workshop on Emerging ICT for Citizen Veillance. Thanks to Emily Goldsher-Diamond for outstanding and invaluable research assistance.

References

Angwin, J., & Valentino-Devries, J. (2012). New tracking frontier: Your license plates. *The Wall Street Journal*. http://online.wsj.com/article/SB10000872396390443995604578004723603576296.html. Accessed June 12, 2014.

Brooks, H. (1980). Technology, evolution, and purpose. *Daedalus*, *109*, 65–81.

[15] Thanks to Ira Rubinstein for suggesting Google Buzz as an illustration of the different thinking generated different interpretations of context. Also, see Ira Rubinstein and Nathan Good (2013).

Cate, F. (2006). The failure of fair information practice principles. In *Consumer protection in the age of the information economy*, July 8. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1156972. Accessed July 1, 2013.

Center for Democracy and Technology. (2012). White House Unveils 'Consumer Privacy Bill of Rights'; Industry Embraces Do Not Track. February 23. https://cdt.org/press/white-house-unveils-consumer-privacy-bill-of-rights-industry-embraces-do-not-track/.

Chavez, P. L. (2011). Comments of Google Inc. to US Department of Commerce. Electronic filing, January 28. http://www.ntia.doc.gov/files/ntia/comments/101214614-0614-01/attachments/FINALCommentsonDepartmentofCommercePrivacyGreenPaper%20%283%29.pdf. Accessed June 11, 2013.

Civil, C. (2012). President Obama's Privacy Bill of Rights: encouraging a collaborative process for digital privacy reform. *Berkeley Technology Law Journal*. http://btlj.org/2012/03/12/president-obamas-privacy-bill-of-rights-encouraging-a-collaborative-process-for-digital-privacy-reform. Accessed June 11, 2013.

Cohen, J. (2012). *Configuring the networked self: Law, code and the play of everyday practice*. New Haven: Yale University Press.

Department of Commerce and National Telecommunications & Information Administration. (2012). Consumer data privacy in a networked world: A framework for protecting privacy and promoting innovation in the global digital economy. White House Privacy Report, February 23. http://www.whitehouse.gov/sites/default/files/privacy-final.pdf. Accessed June 11, 2013.

Department of Homeland Security. (2013). Web site privacy policy. http://www.dhs.gov/privacy-policy. Accessed June 12, 2013.

Dwork, C., & Mulligan, D. K. (2013). It's not privacy, and it's not fair. *Stanford Law Review Online*, *66*, 35.

Electronic Privacy Information Center (2012) White house sets out consumer privacy bill of rights. https://epic.org/2012/02/white-house-sets-out-consumer-.html. Accessed July 9, 2015.

Ellul, J., & Merton, R. K. (1964). *The technological society*. New York: Vintage Books.

European Union. (2013). Committee on Civil Liberties, Justice and Home Affairs. In *On the Proposal for a Regulation of the European Parliament and of the Council on the Protection of Individuals with Regard to the Processing of Personal Data and on the Free Movement of Such Data (General Data Protection Regulation)*. By Jan Phillipp Albrecht. Vol. (COM(2012)0011—C7-0025/2012—2012/0011(COD)).

Federal Trade Commission. (2012). Protecting consumer privacy in an era of rapid change: Recommendations for businesses and policymakers. *FTC Report*. http://www.ftc.gov/os/2012/03/120326privacyreport.pdf. Accessed June 11, 2013.

Federal Trade Commission, Plaintiff, v. Wyndham Worldwide Corporation, Et Al., Defendants. 2:13-cv-01887-ES-JAD. US District Court, District of New Jersey. 7 Apr. 2014.

Friedman, M. (1970). The social responsibility of business is to increase its profits. *The New York Times Magazine*. http://www.colorado.edu/studentgroups/libertarians/issues/friedman-soc-resp-business.html. Accessed June 11, 2013.

Gavison, R. (1980). Privacy and the limits of the law. *Yale Law Journal, 89*, 421–471.

Google Inc v. Joffe Et Al. 13 1181. US Supreme Court. 30 June 2014.

Hoffman, D. (2012). White House releases framework for protecting privacy in a networked world. *Post on Policy@Intel blog*. http://blogs.intel.com/policy/2012/02/23/white-house-privacy. Accessed June 12, 2013.

Horan, P. Re: Information and Privacy in the Internet Economy. Online Publishers Association, January 28. http://www.ntia.doc.gov/files/ntia/comments/101214614-0614-01/attachments/OPA%20Comments%20in%20DOC%20Privacy%20Proceeding%20(Docket%20No.%20101214614-0614-01).pdf. Accessed July 9, 2015.

Intel. (2011). RE: FTC Staff Preliminary Report on Protecting Consumer Privacy. Intel Comments to FTC, January 26. http://www.ftc.gov/os/comments/privacyreportframework/00246-57451.pdf. Accessed June 11, 2013.

Katz, v. United States, (1967), 389 U.S. 347.

Kiseleva, J., Thanh Lam, H., Pechenizkiy, M., & Calders, T. (2013a). Discovering temporal hidden contexts in web sessions for user trail prediction. In *Proceedings of the 22nd international conference on World Wide Web companion* (pp. 1067–1074). International World Wide Web Conferences Steering Committee.

Kiseleva, J., Lam, H. T., Pechenizkiy, M., & Calders, T. (2013b). Predicting Current User Intent with Contextual Markov Models. In *Data mining workshops (ICDMW), 2013 IEEE 13th international conference on* (pp. 391–398). IEEE.

Lawler, B. (2011). Request for comments: Information privacy and innovation in the internet economy. Intuit Comments before the Department of Commerce, Office of the Secretary National Telecommunications and Information Administration, January 28. http://www.ntia.doc.gov/files/ntia/comments/101214614-0614-01/attachments/Intuit.pdf. Accessed June 11, 2013.

Maier, F. (2010). Comments in Response to the Department of Commerce's Green Paper—Commercial Data Privacy & Innovation in the Internet Economy: A Dynamic Policy Framework. Electronic filing, January 28. http://www.ntia.doc.gov/files/ntia/comments/101214614-0614-01/attachments/DoC%20Green%20Paper%20Comments%20(20110128)-Signed.pdf. Accessed 9 July, 2015.

National Telecommunications and Information Administration. (2012). Multistakeholder process to develop consumer data privacy code of conduct concerning mobile application transparency. Notice of meeting published by Federal Register, June 28. https://www.federalregister.gov/articles/2012/06/28/2012-15767/multistakeholder-process-to-develop-consumer-data-privacy-code-of-conduct-concerning-mobile. Accessed June 11, 2013.

Nissenbaum, H. (2010). *Privacy in context: Technology, policy and the integrity of social life*. Stanford, CA: Stanford Law.

Nissenbaum, H. (2011). A contextual approach to privacy online. *Daedalus, 140*(4), 32–48.

Nissenbaum, H. (2012). From preemption to circumvention: If technology regulates why do we need regulation (and Vice Versa)? *Berkeley Technology Law Journal*, *26*, 3.

OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data. (September 23 1980). http://www.oecd.org/internet/ieconomy/oecdguidelinesontheprotectionofprivacyandtransborderflowsofpersonaldata.htm. Accessed on June 11, 2013.

Olmstead, v. United States, (1928), 277 U.S. 438.

Rauhofer, J. (2013). One step forward, two steps back: Critical observations on the proposed reform of the EU data protection framework. *Journal of Law and Economic Regulation, 6*(1).

Raul, A. C., McNicholas, E. R., Brown, C. T., & Adams, J. P. (2011). Comments of AT&T Inc. Before the Department of Commerce Internet Policy Task Force. Federal Trade Commission, January 28. https://www.ftc.gov/sites/default/files/documents/public_comments/preliminary-ftc-staff-report-protecting-consumer-privacy-era-rapid-change-proposed-framework/00420-58060.pdf. Accessed 9 July, 2015.

Re: Netflix Privacy Litigation. No. 11-00379. US District Court, Northern District of California. 6 July 2012. Print.

Regan, P. M. (1995). *Legislating privacy: Technology, social values, and public policy*. Chapel Hill: University of North Carolina Press.

Rubinstein, I. (2010). Privacy and regulatory innovation: Moving beyond voluntary codes. *I/S a Journal of Law and Policy for the Information Society, 6*(3), 356–423.

Rubinstein, I. S., & Good, N. (2013). Privacy by design: A counterfactual analysis of Google and Facebook privacy incidents. *Berkeley Technology Law Journal, 28*, 1333–1583.

Schoeman, F. D. (1984). *Philosophical dimensions of privacy: An anthology*. Cambridge: Cambridge University Press.

Selbst, A. D. (2013). Contextual expectations of privacy. *Cardozo Law Review, 35*, 643–897.

Solove, D. (2006). A taxonomy of privacy. *University of Pennsylvania Law Review, 154*, 477–560.

US Const. amend. VI.

US National Telecommunications and Information Administration. (July, 2013b). *Short form notice.* http://www.ntia.doc.gov/files/ntia/publications/july_25_code_draft.pdf. Accessed June 11, 2013.

US National Telecommunications and Information Administration. Nov. (2013a). *Privacy multistakeholder process: Mobile application transparency—Background*. http://www.ntia.doc.gov/other-publication/2013/privacy-multistakeholder-process-mobile-application-transparency. Accessed June 11, 2013.

USC § 2511(2)(a)(i)—Interception and disclosure of wire, oral, or electronic communications prohibited (2)(a)(i). http://www.gpo.gov/fdsys/granule/USCODE-2011-title18/USCODE-2011-title18-partI-chap119-sec2511/content-detail.html.

Valentino-Devries, J., & Singer-Vine, J. (2012, December 7). They know what you're shopping for. *The Wall Street Journal*. http://www.wsj.com/articles/SB10001424127887324784404578143144132736214.

Van den Hoven, J. M. (1998). Privacy and the varieties of informational wrongdoing. *Austria Journal of Professional and Applied Ethics, 1*(1), 30–43.

Ware, W. H. (1967). *The computer in your future*. Defense Technical Information Center.

World Economic Forum. (2012). Rethinking personal data: Strengthening trust. Report, May. http://www.weforum.org/docs/WEF_IT_RethinkingPersonalData_Report_2012.pdf. Accessed June 11, 2013.