Privacy and Human Behavior In the Age of Information

Alessandro Acquisti, Laura Brandimarte, George Loewenstein[1]

Authors' Pre-Publication, Draft, Non-Final Version

Final version available at http://www.heinz.cmu.edu/~acquisti/cv.htm

and published in Science, 30 January 2015

Vol. 347 no. 6221 pp. 509-514

DOI: 10.1126/science.aaa1465

Abstract. This review summarizes and draws connections between diverse streams of empirical research on privacy behavior. We use three themes to connect insights from social and behavioral sciences: people's *uncertainty* about the consequences of privacy-related behaviors and their own preferences over those consequences; the *context dependence* of people's concern, or lack thereof, about privacy; and the degree to which privacy concerns and *malleable* – manipulable by commercial and governmental interests. Organizing our discussion by these themes, we offer observations concerning the role of public policy in the protection of privacy in the information age.

If this is the age of information, then privacy is the issue of our times. Activities that were once private or shared with the few, now leave trails of data that expose our interests, traits, beliefs, and intentions. We communicate using emails, texts, and social media; find partners on dating sites; learn via online courses; seek responses to mundane and sensitive questions using search engines; read news and books on the cloud; navigate streets with geotracking systems; celebrate our newborns, and mourn our dead, on social media profiles. Through these and other activities, we reveal information - both knowingly and unwittingly - to one another, to commercial entities, and to our governments. The monitoring of personal information is ubiquitous; its storage is so durable as to render one's past undeletable (*1*) – a modern digital skeleton in the closet. Moreover, accompanying this acceleration in data collection, are steady advancements in the ability to aggregate, analyze, and draw sensitive inferences from individuals' data (*2*).

Due to the seismic nature of these developments, there has been considerable debate about their impact on people's welfare; about individuals' ability to navigate a rapidly evolving privacy landscape; and about what, if anything, should be done about privacy at a policy level.

One position in that debate is to embrace the collection of personal data and urge it forward. Both firms and individuals can benefit from the sharing of once hidden data and from applying increasingly sophisticated analytics applied to larger and more interconnected databases (*3*), the argument goes. So too can society as a whole; for instance, electronic medical records can be combined to observe novel drug interactions (*4*). Those who focus on such benefits of data sharing tend to believe that regulatory protection of privacy would only interfere with the fundamentally benign trajectory of information technologies, and with the benefits such technologies may unlock (*5*). Instead of regulation, we should rely on people's ability to make self-interested decisions involving information disclosing and withholding.

An opposing position concentrates on the alarming potential for personal data to be abused – for economic and social discrimination or for hidden influence and manipulation; for coercion, or for censorship. According to this view, the erosion of privacy threatens our autonomy, not merely as consumers but as citizens (*6*); sharing more personal data does not necessarily translate into more

[1] Corresponding author: Alessandro Acquisti, acquisti@andrew.cmu.edu.

progress, efficiency, or equality (*7*). Those who focus on such risks tend to be concerned about the ability of individuals to manage privacy amidst increasingly complex tradeoffs. Traditional tools for privacy decision making, such as choice and consent, no longer provide adequate protection (*8*). Instead of individual responsibility, regulatory intervention may be needed to balance the interests of the subjects of data against the power of commercial entities and governments holding that data.

Are individuals up to the challenge of navigating privacy in the information age? To address this question, we review diverse streams of empirical privacy research from social and behavioral sciences. We highlight research which identifies factors that influence decisions to protect or surrender privacy, and how, in turn, privacy protection or privacy violations affect people's behavior. Information technologies have progressively encroached on every aspect of our personal and professional lives. Thus, the problem of control over personal data has become inextricably linked to problems of personal choice, autonomy, and socio-economic power. Accordingly, this review focuses on the concept of (and literature around) *informational* privacy (that is, privacy of personal data), but it also touches on other conceptions of privacy (such as anonymity or seclusion). Such notions all ultimately relate to the permeable yet pivotal boundaries between public and private (Solove 2006).

We use three themes to organize and draw connections between streams of privacy research that, in many cases, have unfolded independently of each other. The first theme is people's uncertainty about the nature of privacy trade-offs, and their own preferences over them. The second theme is the powerful context-dependence of privacy preferences: The same person can, in some situations, be oblivious to, but in other situations be acutely concerned about, issues of privacy. The third theme is the malleability of privacy preferences, by which we mean that privacy preferences are subject to influence by those possessing greater insight into their determinants. While most individuals are probably unaware of the diverse influences on their concern about privacy, entities whose interests depend on information revelation by others are not. The manipulation of subtle factors that activate or suppress privacy concern can be seen in myriad realms, such as the choice of sharing defaults on social networks, or the provision of greater control on social media, which creates an illusion of safety and encourages greater sharing.

Uncertainty, context dependence, and malleability are closely connected. Context dependence is possible due to uncertainty. Because people are often 'at sea' when it comes to the consequences of, and their feelings about, privacy, they cast around for cues to guide their behavior. Privacy preferences and behaviors are, in turn, malleable and subject to influence, in large part because they are context-dependent and because those with an interest in information divulgence are able to manipulate context to their advantage.

Uncertainty

Individuals manage the boundaries between their private and public spheres in numerous ways: via separateness, reserve, anonymity (*9*), by protecting personal information, but also through deception and dissimulation (*10*). People establish such boundaries for many reasons, including the need for intimacy and psychological respite, and the desire for protection from social influence and control (*11*). Sometimes, these motivations are so visceral and primal that privacy seeking behavior emerges swiftly and naturally. This is often the case when what is intruded upon is one's physical privacy - as when a stranger encroaches in our personal space (*12*,*13,27*) or demonstratively eavesdrops on a conversation. At other times, however, including often when informational privacy is at stake, people experience considerable uncertainty about whether, and to what degree, they should be concerned about privacy.

A first and most obvious source of privacy uncertainty arises from incomplete and asymmetric information. Advancements in information technology have made the collection and usage of personal data often invisible. As a result, individuals rarely have clear knowledge of what information other people, firms, and governments have about them, how that information is used, and with what consequences. To the extent that people lack such information, or are aware of their ignorance, they are likely to be uncertain about how much information to share.

Two factors exacerbate the difficulty of ascertaining the potential consequences of privacy behavior. One factor is that, while some privacy harms are tangible, such as the financial costs associated with identity theft, many others, such as having strangers become aware of one's life history, are intangible. A second factor is that privacy is rarely an unalloyed good – it typically involves tradeoffs (14). For example, ensuring the privacy of a consumer's purchases can protect one from price discrimination, but also deny one the potential benefits of targeted advertisements and offers.

Elements that mitigate one or both of these exacerbating factors, by either increasing the tangibility of privacy harms or making tradeoffs explicit and simple to understand, will generally affect privacy-related decisions. This is illustrated by one lab experiment in which participants were asked to use a specially designed search engine to find online merchants and purchase from them, with their own credit cards, either a set of batteries or a sex toy (*15*). When the search engine only provided links to the merchants' sites and a comparison of the products' prices from the different sellers, a majority of participants did not pay any attention to the merchants' privacy policies; they purchased from those offering the lowest price. However, when the search engine also provided participants with salient, easily accessible, information about the differences in privacy protection afforded by the various merchants, a majority of participants paid a roughly 5% premium to buy products from (and share their credit card information with) more privacy-protecting merchants.

A second source of privacy uncertainty relates to preferences. Even when aware of the consequences of privacy decisions, people are still likely to be uncertain about their own privacy preferences. Research on preference uncertainty (*16*) shows that individuals often have little sense of how much they like goods, services, or other people. Privacy does not seem to be an exception. This can be illustrated by research in which people were asked sensitive and potentially incriminating questions either point-blank, or followed by credible assurances of confidentiality (*17*). Although logically such assurances should lead to greater divulgence, they often had the opposite effect, because they elevate respondents' privacy concerns which, without assurances, would have remained dormant.

The remarkable uncertainty of privacy preferences comes into play in efforts to measure individual differences in preference for privacy (*18*). For example, Alan Westin (*19*) famously used broad (that is, not contextually specific) privacy questions in surveys to cluster individuals into privacy segments: privacy fundamentalists, pragmatists, and unconcerned. When asked directly, many people fall in the first segment: they profess to care a lot, and express particular concern over losing control of their personal information or unauthorized access to it (*20-21*). However, doubts about the power of attitudinal scales to predict actual privacy behavior arose early in the literature (*22*). This discrepancy between attitudes and behaviors has become known as the 'privacy paradox.'

In one early study illustrating the paradox, participants were first classified into categories of privacy concern, inspired by Westin's categorization, based on their responses to a survey dealing with attitudes towards sharing data (*23*). Next, participants were presented with products to purchase at a discount with the assistance of an anthropomorphic shopping agent. Few, regardless of the group they were categorized in, exhibited much reluctance to answer the increasingly sensitive questions the agent plied them with.

Why do people who claim to care about privacy often show little concern about it in their daily behavior? One possibility is that the paradox is illusory – that privacy attitudes (which are defined broadly), and intentions, and behaviors (which are defined narrowly) should not be expected to be closely related (*24, 25*). Thus, one might care deeply about privacy in general, but, depending on the costs and benefits (*26*) prevailing in a specific situation, seek or not seek privacy protection.

This explanation for the privacy paradox, however, is not entirely satisfactory for two reasons. The first one is that the explanation does not account for situations in which attitude-behavior dichotomies arise under high correspondence between expressed concerns and behavioral actions. Consider, for instance, a study that compared attitudinal survey answers to actual social media behavior (*29*). Within the subset of participants who expressed the highest degree of concern over strangers being able to easily find out their sexual orientation, political views, and partners' names, 48% did in fact publicly reveal their sexual orientation online, 47% their political orientation, and 21% their current partner's name. The second reason is that privacy decision making is only in part the result of a rational "calculus" of costs and benefits (14, 26); it is also affected by misperceptions of those costs and benefits, as well as social norms, emotions, and heuristics. Any of these factors may affect behavior differently from how they affect attitudes. For instance, present-bias can cause even the privacy-conscious to engage in risky revelations of information, if the immediate gratification from disclosure trumps the delayed, and hence discounted, future consequences (*30*).

Preference uncertainty is evident not only in studies that compare stated attitudes with behaviors, but also in those that estimate monetary valuations of privacy. 'Explicit' investigations ask people to make direct trade-offs, typically between privacy of data and money. For instance, in a study conducted both in Singapore and the U.S., students made a series of hypothetical choices about sharing information with websites that differed in protection of personal information and prices for accessing services (31). Using conjoint analysis, the authors concluded that subjects valued protection against errors, improper access, and secondary use of personal information between $30.49 and $44.62. Like direct questions about attitudes and intentions, such explicit investigations of privacy valuation spotlight privacy as an issue that respondents should take account of, and, as a result, increase the weight they place on privacy in their responses.

Implicit investigations, in contrast, infer valuations of privacy from day-to-day decisions in which privacy is only one of many considerations, and is typically not highlighted. Individuals engage in privacy related transactions all the time, even when the privacy trade-offs may be intangible, or when the exchange of personal data may not be a visible or primary component of a transaction. For instance, completing a query on a search engine is akin to selling personal data (one's preferences and contextual interests) to the engine in exchange for a service (search results). "Revealed preference" economic arguments would then conclude that, since technologies for information sharing have been enormously successful, while technologies for information protection have not, individuals hold overall low valuations of privacy. However, that is not always the case: while individuals at times give up personal data for small benefits or discounts, at other times they can voluntarily incur substantial costs to protect their privacy. Context, as further discussed in the next section, matters.

In fact, attempts to pinpoint exact valuations that people assign to privacy may be misguided, as suggested by research calling into question the stability, and hence validity, of privacy estimates. In one field experiment inspired by the literature on endowment effects (*32*), shoppers at a mall were offered gift cards for participating in a non-sensitive survey. Participants had to choose between a $10 "anonymous" gift card (transactions done with that card would not be traceable to the subject) or a $12 trackable card (transactions done with that card would be linked to the name of the subject). Both cards could be used online or in stores, just like debit cards. Initially, half of the participants were given one type of card, and half with the other. Then, they were all offered the opportunity to switch. Some shoppers, for example, were given the anonymous $10 card and were asked if they would accept $2 to "allow my name to be linked to transactions done with the card;" for other subjects, the question was whether they would accept a card with $2 less value to "prevent my name from being linked to transactions done with the card." Five times as many subjects (52.1%) originally holding the less valuable but anonymous card, chose it over the other card, than those originally holding the more valuable card (9.7%). This suggests that people value privacy more when they have it than when they do not.

The consistency of preferences for privacy is also complicated by the existence of a powerful countervailing motivation: the desire to be public, share, and disclose. Humans are social animals, and information sharing is a central feature of human connection. Social penetration theory (*33*) suggests that progressively increasing levels of self-disclosure are an essential feature of the natural and desirable evolution of interpersonal relationships from superficial to intimate. Such a progression is only possible when people begin social interactions with a baseline level of privacy. Paradoxically, therefore, privacy provides an essential foundation for intimate disclosure. Like privacy, self-disclosure confers numerous objective and subjective benefits, including psychological and physical health (*34-35*). The desire for interaction, socialization, disclosure, and recognition or fame (and, conversely, the fear of anonymous insignificance) are fundamental human motives, like the need for privacy. The electronic media of the current age provide unprecedented opportunities for acting on them. Through social media, disclosures can build social capital or increase self-esteem (*36*), and fulfill ego needs (*37*). In a series of fMRI experiments, self-disclosure was even found to engage neural mechanisms associated with reward: people highly value the ability to share thoughts and feelings with others. So much so, subjects in one of the experiments were even willing to forgo money in order to disclose about themselves (*38*).

Context Dependence

Much evidence suggests that privacy is a universal human need (*39*, and Box 1). However, when people are uncertain about their preferences, they often search for cues in their environment to provide guidance. And, since cues are a function of context, behavior is as well. Applied to privacy, context-dependence means that individuals can, depending on the situation, exhibit anything ranging from extreme concern to apathy about privacy. Adopting the terminology of Westin, we are all privacy pragmatists, fundamentalists, or unconcerned, depending on time and place (*28*).

The way we construe and negotiate public and private spheres is context-dependent because the boundaries between the two are murky (*40*): The rules people follow for managing privacy vary by situation, are learned over time, and are based on cultural, motivational, and purely situational criteria. For instance, we may be more comfortable sharing secrets with friends, but at times be caught by surprise by our revelations to a stranger on a plane (*41*). The theory of contextual 'integrity' posits that social expectations affect our beliefs regarding what is private and what is public, and that such expectations vary with specific contexts. (*42*). Thus, depending on context, seeking privacy in public is not an oxymoron: individuals can manage privacy even while sharing information, and even on social media (boyd 2014). For instance, a longitudinal study of actual disclosure behavior of online social network users highlighted that, over time, many users increased the amount of personal information revealed to their friends (those connected to them on the network), while simultaneously decreasing the amounts revealed to strangers (those unconnected to them) (*43* and Figure 1).

The cues that people use to judge the importance of privacy sometimes result in sensible behavior. For instance, the presence of government regulation has been shown to reduce consumer concern and increase trust: it is a cue that people use to infer that some degree of privacy protection will be afforded (Xu et al 2009). In other situations, though, cues can be unrelated, or even negatively related, to normative bases of decision making. For example, in one online experiment (*45*), individuals were more likely to reveal personal and even incriminating information on a website with an unprofessional and cheesy design with the banner "How Bad R U" than on a site with a formal interface – even though the site with the formal interface was judged by other respondents to be much safer (Figure 2). Yet in other situations, it is the physical environment that influences privacy concern and associated behavior *(46)*, sometimes even unconsciously. For instance, all else being equal, intimacy of self-disclosure is higher in warm, comfortable rooms, with soft lighting, than in cold, non-intimate rooms, with bare cement and over-head fluorescent lighting (*47*).

Some of the cues that influence perceptions of privacy are one's culture, and the behavior of other people - either through the mechanism of descriptive norms (imitation) or via reciprocity (*49*). Observing other people reveal information increases the likelihood one will reveal it oneself (*48*). In one study, survey-takers were asked a series of sensitive personal questions regarding their engagement in illegal or ethically questionable behaviors. After answering each question, participants were provided with information about the share of other participants who, in the same survey, had admitted to having engaged in a given behavior. Unbeknownst to them, that information had been manipulated. Being provided with information which suggested that a majority of survey takers had admitted a certain questionable behavior increased participants' willingness to disclose their engagement in *other*, also sensitive, behaviors. Other studies find that the tendency to reciprocate information disclosure is so ingrained that people will reveal more information even to a computer agent that provides information about itself (*50*). Findings such as this may help to explain the escalating amounts of self-disclosure we witness online: if others are doing it, people seem to reason unconsciously, doing so oneself must be desirable or safe.

Other people's behavior affects privacy concerns in other ways too. Sharing personal information with others makes them "co-owners" of that information (*51*), and, as such, responsible for its protection. Mismanagement of shared information by one or more co-owners causes "turbulence" of the privacy boundaries and, consequently, negative reactions, including anger or mistrust. In a study of undergraduate Facebook users (*52*), for instance, turbulence of privacy boundaries, due to having one's profile exposed to unintended audiences, dramatically increased the odds that a user would restrict profile visibility to friends only.

Likewise, privacy concerns are often a function of past experiences. When something in an environment changes, like the introduction of a camera or other monitoring devices, privacy concern is likely to be activated. For instance, surveillance can produce discomfort (*53*) and negatively affect worker productivity (*54*). However, privacy concern, like other motivations, is adaptive; people get used to levels of intrusion that are unchanging over time. In an experiment conducted in Helsinki (*55*), the installation of sensing and monitoring technology in households led family members to change their behavior initially, particularly in relation to conversations, nudity, and sex. And yet, if they accidentally performed an activity, such as walking naked into the kitchen in front of the sensors, it seemed to have the effect of 'breaking the ice'; participants then showed less concern about repeating the behavior. More generally, participants became inured to the presence of the technology over time.

The context-dependence of privacy concern has significant implications for the risks associated with modern information and communication technology (*56*). With online interactions we no longer have a clear sense of the spatial boundaries of our listeners. Who is reading our blog post? Who is looking at our photos online? Adding complexity to privacy decision making, boundaries between public and private become even less defined in the online world (*57*) where we become social media friends with our co-workers and post pictures to an indistinct flock of followers. With different social groups mixing online, separating online and offline identities and meeting our and others' expectations regarding privacy, information sharing becomes harder (*58*).

Malleability and Influence

While individuals are often unaware of the diverse factors that determine their concern about privacy in a particular situation, entities whose prosperity depends on information revelation by others are much more sophisticated. With the emergence of the information age, growing institutional and economic interests have developed around disclosure of personal information, from online social networks to behavioral advertising. It is not surprising, therefore, that some entities have an interest in, and have developed expertise in, exploiting behavioral and psychological processes to promote disclosure (*59*). Such efforts play on the malleability of privacy preferences, a term we use to refer to the observation that various, sometimes subtle, factors can be used to activate or suppress privacy concerns, which in turn affect behavior.

Defaults are an important tool used by different entities to affect information disclosure. A large body of research has shown that defaults matter for decisions as important as organ donation and retirement saving (*60*). Sticking to defaults is convenient, and people often interpret defaults as implicit recommendations (*61*). Thus, it is not surprising that default settings for one's profile's visibility on social networks (*62*), or the existence of opt-in or opt-out privacy policies on websites (*63*) affect individuals' privacy behavior (Figure 3).

In addition to defaults, websites can also employ design features that frustrate or even confuse users into disclosing personal information (*64*), a practice that has been referred to as "malicious interface design" (*65*). Another obvious strategy that commercial entities can employ to avoid raising privacy concerns is to not 'ring alarm bells' when it comes to data collection. When companies do ring them, for example by using overly fine-tuned personalized advertisements, consumers are alerted (*66*) and can respond with negative "reactance" (*67*).

Various so-called "antecedents" (*68*) affect privacy concerns, and can be used to influence privacy behavior. For instance, trust in the entity receiving one's personal data soothes concerns. Moreover, because some interventions intended to protect privacy can inculcate trust, concerns can be muted by the very interventions intended to protect privacy. Perversely, 62% of respondents to a survey believed (incorrectly) that the existence of a privacy policy implied that a site could not share their personal information without permission (*28*), which suggests that simply posting a policy that consumers don't read may lead to misplaced feelings of being protected.

Control is another feature that can inculcate trust and produce paradoxical effects. Perhaps due to its lack of controversiality, control has been one of the capstones of the focus of both industry and policy makers in attempts to balance privacy needs against the value of sharing. Control over personal information is often perceived as a critical feature of privacy protection (*39*). In principle, it does provide users with the means to manage access to their personal information. Research, however, shows that control can reduce privacy concern (*69*), which in turn can have unintended effects. For instance, one study we conducted found that users of a social website who were provided with greater explicit control over whether and what information to publish ended up sharing more information –the opposite of the ostensible purpose of providing such control (*70*).

Similar to the normative perspective on control, increasing the transparency of firms' data practices would seem to be desirable, for instance via the simplification and more explicit posting of privacy policies. However, transparency mechanisms can be easily rendered ineffective. Research has highlighted not only that an overwhelming majority of Internet users do not read privacy policies (*71*), but also that few users would benefit from doing so: nearly half of a sample of online privacy policies were found to be written in language beyond the grasp of most Internet users (*72*). Indeed, and somewhat amusingly, it has been estimated that the aggregate opportunity cost of US consumers actually reading privacy policies of the sites they visit would be $781 billion/year (*73*).

Although uncertainty and context-dependence lead naturally to malleability and manipulation, not all malleability is necessarily sinister. Consider monitoring. While monitoring can cause discomfort and reduce productivity, the feeling of being observed and accountable can incentivize people to engage in pro-social behaviors, or (for better or for worse) adhere to social norms (*74*). Pro-social behavior can even be heightened by monitoring cues such as three dots in a stylized face configuration (*75*). By the same token, the depersonalization induced by computer-mediated interaction (*76*), either in the form of lack of identifiability or of visual anonymity (*77*), can have beneficial effects, such as increasing truthful responses to sensitive surveys (*78-79*). Whether elevating or suppressing privacy concerns is socially beneficial critically depends, yet again, on context (for a meta-analysis of the impact of de-identification on behavior, see *80*). For example, perceptions of anonymity can alternatively lead to dishonest or prosocial behavior. Illusory anonymity induced by darkness caused participants in an experiment (*81*) to cheat in order to gain more money. This can be interpreted as a form of disinhibition effect (*82*), by which perceived anonymity licenses people to act in ways that they would otherwise not even consider. In other circumstances, though, anonymity leads to prosocial behavior - for instance higher willingness to share money in a dictator game, when coupled with priming of religiosity (*83*).

Conclusions

Norms and behaviors regarding private and public greatly differ across cultures (8*4*). Americans, for example, are reputed to be more open about sexual matters than the Chinese, while the latter more open about financial matters (e.g., income, cost of home, and possessions). And even within cultures, people differ substantially in how much they care about privacy and what information they treat as private. More importantly, as we have sought to highlight in this review, privacy concerns can vary dramatically for the same individual, and for societies, over time.

If privacy behaviors are culture and context dependent, however, the dilemma of what to share and what to keep private is universal across societies and over human history. The task of navigating those boundaries, and the consequences of mismanaging them, have grown increasingly complex and fateful in the information age, to the point where our natural instincts seem not nearly adequate.

In this review we used three themes to organize and draw connections between the social and behavioral science literatures on privacy and behavior. We end the review with a brief discussion of the reviewed literature's relevance to privacy policy.

Uncertainty and context dependence imply that people cannot always be counted upon to navigate the complex tradeoffs involving privacy in a self-interested fashion. People are often unaware of the information they are sharing, unaware of how it can be used, and even in the rare situations when they have full knowledge of the consequences of sharing, uncertain about their own preferences. Malleability, in turn, implies that people are easily influenced in what and how much they disclose. Moreover, what they share can be used to influence their emotions, thoughts, and behaviors in many aspects of their lives: as individuals, as consumers, and as citizens. While such influence is not always or necessarily malevolent or dangerous, relinquishing control over one's personal data and over one's privacy implies alters the balance of power between those holding the data and those who are the subjects of that data.

Insights from the social and behavioral empirical research on privacy we have reviewed suggests that policy approaches that rely exclusively on informing or “empowering” the individual are likely to provide inadequate protection against the risks posed by novel information technologies. Consider transparency and control, two principles conceived as necessary conditions for privacy protection. The research we highlighted shows that they provide radically insufficient protections, and may even backfire when used apart from other principles of data protection. Specifically, the research reviewed here suggests that, *if* the goal of policy is to adequately protect privacy (as we believe it should be), then we need policies that protect individuals with minimal requirement of informed and rational decision making, and policies that include a baseline framework of protection - such as the principles embedded in so-called fair information practices (*86*). People need assistance and even protection to aid in balancing what is otherwise a very uneven playing field. As highlighted by our discussion, a goal of policy should be to achieve a more even equity of power between individuals, consumers, and citizens on the one hand, and, on the other, the data holders such as governments and corporations that currently have the upper hand. To be effective, privacy policy should protect the naïve, the uncertain, and the vulnerable. It should be sufficiently flexible to evolve with the emerging, unpredictable, complexities of the information age.

References

1. V. Mayer-Schönberger, *Delete: The Virtue of Forgetting In the Digital Age* (Princeton Univ. Press, Princeton, 2011).
2. L. Sweeney, K-anonymity: A model for protecting privacy. *Int. J. Uncertainty, Fuzziness & Knowl.-Based Syst*. **10**, 557-570 (2002).
3. A. McAfee, E. Brynjolfsson, Big data. The management evolution. *Harvard Bus. Rev.* **90**, 61-67 (2012).
4. N.P. Tatonetti, P.P. Ye, R. Daneshjou, R.B. Altman, Data-driven prediction of drug effects and interactions. *Sci. Transl. Med.* **4**, 125ra31 (2012).
5. R.A. Posner, The economics of privacy. *Am. Econ. Rev*. **71**, 405-409 (1981).
6. J.E. Cohen, Examined lives: informational privacy and the subject as object*. Stanford L. Rev.* **52**, 1373-1438 (2000).
7. K. Crawford, K. Miltner, M.L. Gray, Critiquing big data: politics, ethics, epistemology. *Int. J. Comm.* **8**, 1663-1672 (2014).
8. D.J. Solove, Introduction: Privacy self-management and the consent dilemma. *Harvard L. Rev.* **126**, 1880-1903 (2013).
9. F. Schoeman, Ed., *Philosophical dimensions of privacy – An anthology* (Cambridge University Press, New York, 1984).
10. B.M. DePaulo, C. Wetzel, W. Sternglanz, M. J. W. Wilson, Verbal and nonverbal dynamics of privacy, secrecy, and deceit. *J. Soc. Iss.* **59**, 391-410 (2003).
11. S.T. Margulis, Privacy as a social issue and behavioral concept. *J. Soc. Iss.* **59**, 243-261 (2003).
12. E. Goffman, *Relations in Public: Microstudies of the Public Order* (Harper & Row, New York, 1971.
13. E. Sundstrom, I. Altman, Interpersonal relationships and personal space: Research review and theoretical model. *Hum. Ecol.* **4**, 47-67 (1976).
14. R.S. Laufer, M. Wolfe, Privacy as a concept and a social issue: a multidimensional developmental theory. *J. Soc. Issues* **33**, 22-42, (1977).
15. J.Y. Tsai, S. Egelman, L. Cranor, A. Acquisti, The effect of online privacy information on purchasing behavior: an experimental study. *Inf. Syst. Res.* **22**, 254-268 (2011).
16. P. Slovic, The construction of preference. *Am. Psychol*. **50**, 364-371 (1995).
17. E. Singer, H. Hippler, N. Schwarz, Confidentiality assurances in surveys: reassurance or threat? *Int. J. Public Opin. Res.* **4**, 256-268 (1992).
18. V.P. Skotko, D. Langmeyer, The effects of interaction distance and gender on self-disclosure in the dyad. *Sociometry* **40**, 178–182 (1977).
19. A. Westin, Harris Louis & Associates, Harris-Equifax Consumer Privacy Survey. (Tech. rep. 1991).
20. M.J. Culnan, P. K. Armstrong, Information privacy concerns, procedural fairness, and impersonal trust: an empirical investigation. *Organization Sci*. **10**, 104-115 (1999).
21. H.J. Smith, S.J. Milberg, S.J. Burke, Information privacy: Measuring individuals' concerns about organizational practices. *MIS Quarterly* **20**, 167-196 (1996).
22. B. Lubin, R.L. Harrison, Predicting small group behavior with the self-disclosure inventory. *Psychol. Rep.* **15**, 77-78 (1964).

23. S. Spiekermann, J. Grossklags, B. Berendt, *E-Privacy in 2nd Generation E-Commerce: Privacy Preferences Versus Actual Behavior* (3rd ACM Conference on Electronic Commerce, Tampa, 2001), pp. 38-47.
24. P.A. Norberg, D.R. Horne, D.A. Horne, The privacy paradox: Personal information disclosure intentions versus behaviors. *J. Consumer Affairs* **41**, 100-126 (2007).
25. I. Ajzen, M. Fishbein, Attitude-behavior relations: a theoretical analysis and review of empirical research. *Psychol. Bull.* **84**, 888-918 (1977).
26. P.H. Klopfer, D.I. Rubenstein, The concept privacy and its biological basis. *J. Soc. Iss.* **33**, 52-65 (1977).
27. B. Schwartz, The social psychology of privacy. *Am. J. Sociology* **73**, 741-752 (1968).
28. C.J. Hoofnagle, J.M. Urban, Alan Westin's privacy homo economicus. *Wake Forest L. Rev.* **49**, 261-321 (2014).
29. A. Acquisti, R. Gross, in *Privacy enhancing technologies*, G. Danezis & P. Golle Eds. (Springer, New York, 2006), pp. 36-58.
30. A. Acquisti, *Privacy in Electronic Commerce and the Economics of Immediate Gratification* (5th ACM Conference on Electronic Commerce, New York, 2004), pp 21-29.
31. I. Hann, K. Hui, S.T. Lee, I.P.L. Png, Overcoming online information privacy concerns: an information-processing theory approach. *J. of Mgmt. Inf. Syst*. **24**, 13-42 (2007).
32. A. Acquisti, L.K. John, G. Loewenstein, What is privacy worth? *J. Leg. St.* **42**, 249-274 (2013).
33. I. Altman, D. Taylor, *Social Penetration: The Development of Interpersonal Relationships* (Holt, Rinehart & Winston, NY, 1973).
34. J. Frattaroli, Experimental disclosure and its moderators: a meta-analysis. *Psychol. Bull.* **136**, 823-865 (2006).
35. J.W. Pennebaker, Putting stress into words: health, linguistic, and therapeutic implications. *Behav. Res. Ther*. **31**, 539-548 (1993).
36. C. Steinfield, N.B. Ellison, C. Lampe, Social capital, self-esteem, and use of online social network sites: A longitudinal analysis. *J. Appl. Dev. Psychol.* **29**, 434-445 (2008).
37. C.L. Toma, J.T. Hancock, Self-affirmation underlies Facebook use. *Pers. Soc. Psychol. Bull.* **39**, 321-331 (2013).
38. D.I. Tamir, J.P. Mitchell, Disclosing information about the self is intrinsically rewarding. *PNAS* **109**, 8038–8043 (2012).
39. A. Westin, *Privacy and Freedom* (Athenäum, New York, 1967).
40. G. Marx, Murky conceptual waters: The public and the private. *Ethics & Inf. Tech.* **3**, 157-169 (2001).
41. J.W. Thibaut, H.H. Kelley, *The Social Psychology Of Groups* (Wiley, Oxford, 1959).
42. H. Nissenbaum, *Privacy in Context: Technology, Policy, and the Integrity of Social Life* (Stanford Univ. Press, Redwood City, 2009).
43. D. Boyd, It’s Complicated: The Social Lives of Networked Teens (Yale Univ. Press, New Haven, 2014)
44. F. Stutzman, R. Gross, A. Acquisti, Silent listeners: The evolution of privacy and disclosure on facebook. *J. Privacy & Confidentiality* **4**, 7-41 (2013).
45. D.J. Solove, A taxonomy of privacy. *Univ. Penn. L. Rev.* **154**, 477-564 (2006).

46. L.K. John, A. Acquisti, G. Loewenstein, Strangers on a plane: Context-dependent willingness to divulge sensitive information. *J. Consumer Res.* **37**, 858-873 (2011).
47. I. Altman, *The Environment and Social Behavior: Privacy, Personal Space, Territory, and Crowding* (Cole, Monterey, 1975).
48. A.L. Chaikin, V.J. Derlega, S.J. Miller, Effects of room environment on self-disclosure in a counseling analogue. *J. Counseling Psychol*. **23**, 479-481 (1976).
49. A. Acquisti, L.K. John, G. Loewenstein, The impact of relative standards on the propensity to disclose. *J. Marketing Res.* **49**, 160-174 (2012).
50. V.J. Derlega, A.L. Chaikin, Privacy and self-disclosure in social relationships. *J. Soc. Iss.* **33**, 102-115 (1977).
51. Y. Moon, Intimate exchanges: Using computers to elicit self-disclosure from consumers. *J. Consumer Res*. **26**, 323-339 (2000).
52. S. Petronio, *Boundaries of Privacy: Dialectics of Disclosure* (SUNY Press, Albany, 2002).
53. F. Stutzman, J. Kramer-Duffield. *Friends Only: Examining a Privacy-Enhancing Behavior in Facebook* (SIGCHI Conference on Human Factors in Computing Systems, ACM, Atlanta, 2010) pp. 1553-1562.
54. T. Honess, E. Charman, *Closed Circuit Television in Public Places: Its Acceptability and Perceived Effectiveness*, (Police Research Group, London, 1992).
55. M. Gagné, E.L. Deci, Self-determination theory and work motivation. *J. Organizational Behav*. **26**, 331-362 (2005).
56. A. Oulasvirta *et al*., *Long-Term Effects of Ubiquitous Surveillance in the Home* (ACM Conference on Ubiquitous Computing, Pittsburgh, 2012), pp 41-50.
57. L. Palen, P. Dourish, *Unpacking “Privacy” For A Networked World* (SIGCHI Conference on Human Factors in Computing Systems, ACM, Fort Lauderdale, 2003), pp. 129-136.
58. Z. Tufekci, Can you see me now? Audience and disclosure regulation in online social network sites. *Bull. Sci. Tech. Soc.* **28**, 20-36 (2008).
59. J.A. Bargh, K.Y.A. McKenna, G.M. Fitzsimons, Can you see the real me? Activation and expression of the “true self” on the internet. *J. Soc. Iss.* **58**, 33-48 (2002).
60. R. Calo, Digital market manipulation. *Geo. Wash. L. Rev.* **82**, 995-1304 (2014).
61. E.J. Johnson, D. Goldstein, Do defaults save lives? *Sci.* **302**, 1338–1339 (2003).
62. C.R. McKenzie, M.J. Liersch, S.K. Finkelstein, Recommendations implicit in policy defaults. *Psychol. Sci.* **17**, 414-20 (2006).
63. R. Gross, A. Acquisti, *Information Revelation and Privacy in Online Social Networks* (ACM Workshop – Privacy in the Electronic Society, New York, 2005), pp.71-80.
64. E.J. Johnson, S. Bellman, G.L. Lohse, Defaults, framing and privacy: Why opting in ≠ opting out, *Marketing Letters* **13**, 5-15 (2002).
65. W. Hartzog, Website design as contract. *Am. U. L. Rev.* **60**, 1635-1671 (2010).
66. G. Conti, E. Sobiesk, *Malicious Interface Design: Exploiting the user* (19th International Conference on World Wide Web, ACM, Raleigh, 2010) pp. 271-280.
67. A. Goldfarb, C. Tucker, Online display advertising: Targeting and obtrusiveness. *Marketing Sci.* **30**, 389-404 (2011).
68. T.B. White, D.L. Zahay, H. Thorbjørnsen, S. Shavitt, Getting too personal: Reactance to highly personalized email solicitations. *Marketing Letters*, **19**, 39-50 (2008).

69. H.J. Smith, T. Dinev, H. Xu, Information privacy research: An interdisciplinary review. *MIS Quarterly* **35**, 989-1016 (2011).
70. H. Xu, H.H. Teo, B.C. Tan, R. Agarwal, The role of push-pull technology in privacy calculus: The case of location-based services. *J. of Mgmt Inf. Syst.* **26**, 135-174 (2009).
71. L. Brandimarte, A. Acquisti, G. Loewenstein, Misplaced confidences privacy and the control paradox. *Soc. Psychol. Pers. Sci.* **4**, 340-347 (2013).
72. C. Jensen, C. Potts, C. Jensen, Privacy practices of Internet users: Self-reports versus observed behavior. *Int. J. Human-Computer Stud.* **63**, 203-227 (2005).
73. C. Jensen, C. Potts, *Privacy Policies as Decision-Making Tools: An Evaluation of Online Privacy Notices* (SIGCHI Conference on Human factors in computing systems, ACM, Vienna, 2004), pp. 471-478.
74. A.M. McDonald, L.F. Cranor, The cost of reading privacy policies, *I/S: J. L. Policy Inf. Society.* **4**, 540-565 (2008).
75. C. Wedekind, M. Milinski, Cooperation through image scoring in humans. *Sci*. **288**, 850-852 (2000).
76. M. Rigdon, K. Ishii, M. Watabe, S. Kitayama, Minimal social cues in the dictator game. *J. Econ. Psychol.* **30**, 358-367 (2009).
77. S. Kiesler, J. Siegel, T.W. McGuire, Social psychological aspects of computer-mediated communication. *Am. Psychol.* **39**, 1123-1134 (1984).
78. A.N. Joinson, Self-disclosure in computer-mediated communication: the role of self-awareness and visual anonymity. *Eu. J. Soc. Psychol.* **31**, 177-192 (2001).
79. S. Weisband, S. Kiesler, *Self Disclosure On Computer Forms: Meta-Analysis And Implications* (SIGCHI Conference Conference on Human Factors in Computing Systems, ACM, Vancouver, 1996), pp 3-10.
80. R. Tourangeau, T. Yan, Sensitive questions in surveys. *Psychol. Bull.* **133**, 859-883 (2007).
81. T. Postmes, R. Spears, Deindividuation and antinormative behavior: A meta-analysis. *Psychol. Bull.* **123**, 238-259 (1998).
82. C. Zhong, V.K. Bohns, F. Gino, Good lamps are the best police: Darkness increases dishonesty and self-interested behavior. *Psychol. Sci.* **21**, 311-314 (2010).
83. J. Suler, The online disinhibition effect. *Cyberpsychol. & Behav.* **7**, 321-326 (2004).
84. A.F. Shariff, A. Norenzayan**,** God is watching you: Priming God concepts increases prosocial behavior in an anonymous economic game. *Psychol. Sci.* **18**, 803-809 (2007).
85. B. Moore, *Privacy: Studies in Social and Cultural History* (Armonk, NY, 1984).
86. C. Camerer *et al.*, Regulation for Conservatives: Behavioral Economics and the Case for Asymmetric Paternalism. *Univ. Penn. L. Rev.* **151**, 1211-1254 (2003).
87. *Records, Computers and the Rights of Citizens* (Secretary's Advisory Committee, U.S. Dept. of Health, Education and Welfare, Washington, DC, 1973).

Endogenous privacy behavior and exogenous shocks. Percentage of profiles on the Carnegie Mellon University Facebook network who revealed birthday and high school over time, 2005-2011.

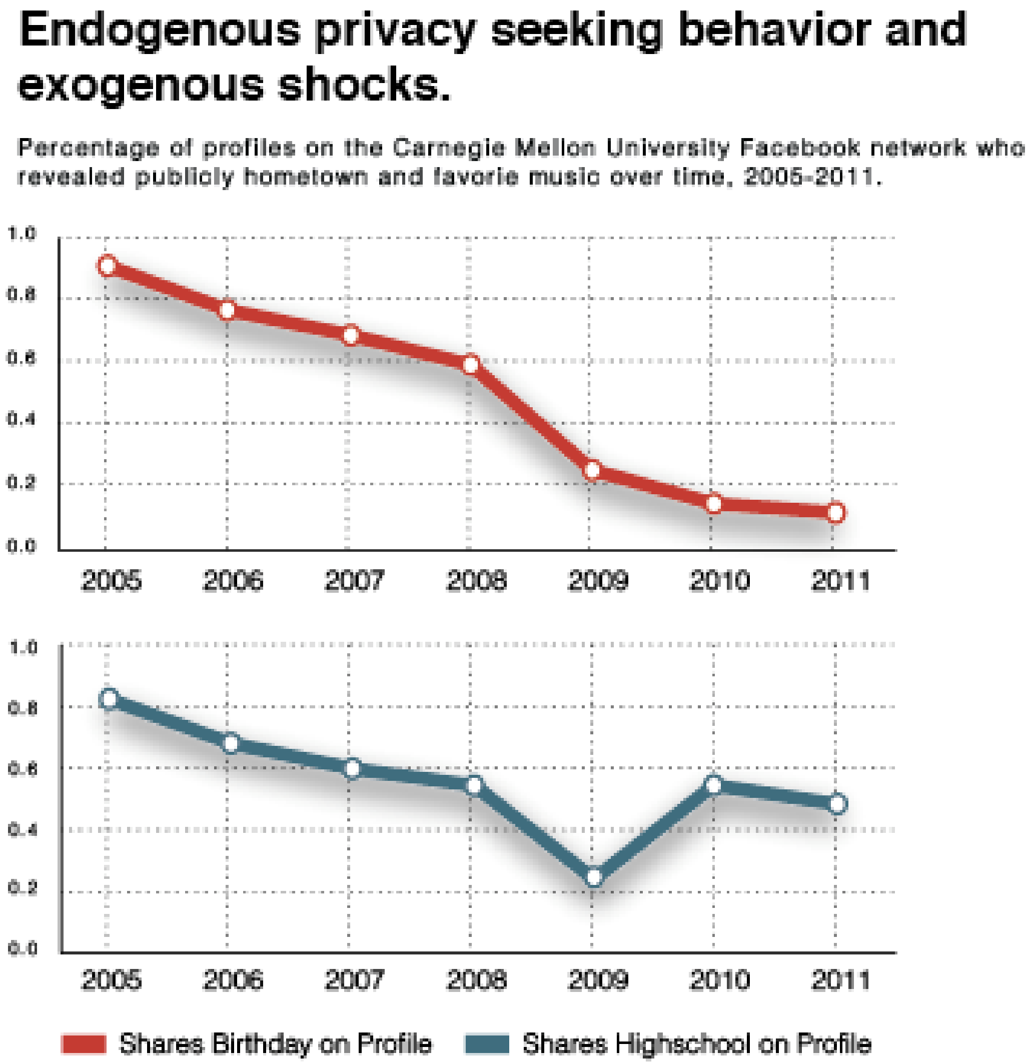

Figure 1. Endogenous privacy behavior and exogenous shocks. Privacy behavior is affected both by endogenous motivations (for instance, subjective preferences) and exogenous factors (for instance, changes in user interfaces). Over time, the percentage of members in the Carnegie Mellon University Facebook network who chose to publicly reveal personal information decreased dramatically. For instance, over 80% of profiles revealed publicly their birthday in 2005, but fewer than 20% in 2011, suggestive of an increase in privacy seeking over time. The decreasing trend is not uniform, however: for instance, after decreasing for several years, suddenly the percentage of profiles that publicly revealed their high school roughly doubled between 2009 and 2010 – after Facebook changed the default visibillity settings for various fields on its profiles, including high school (bottom), but not Birthday (top) (43).

The impact of cues on disclosure behavior. Relative admission rates, by experimental condition, in an experiment testing the impact of different survey interfaces on willigness to answer questions about the subject's engagement in various sensitive behaviors.

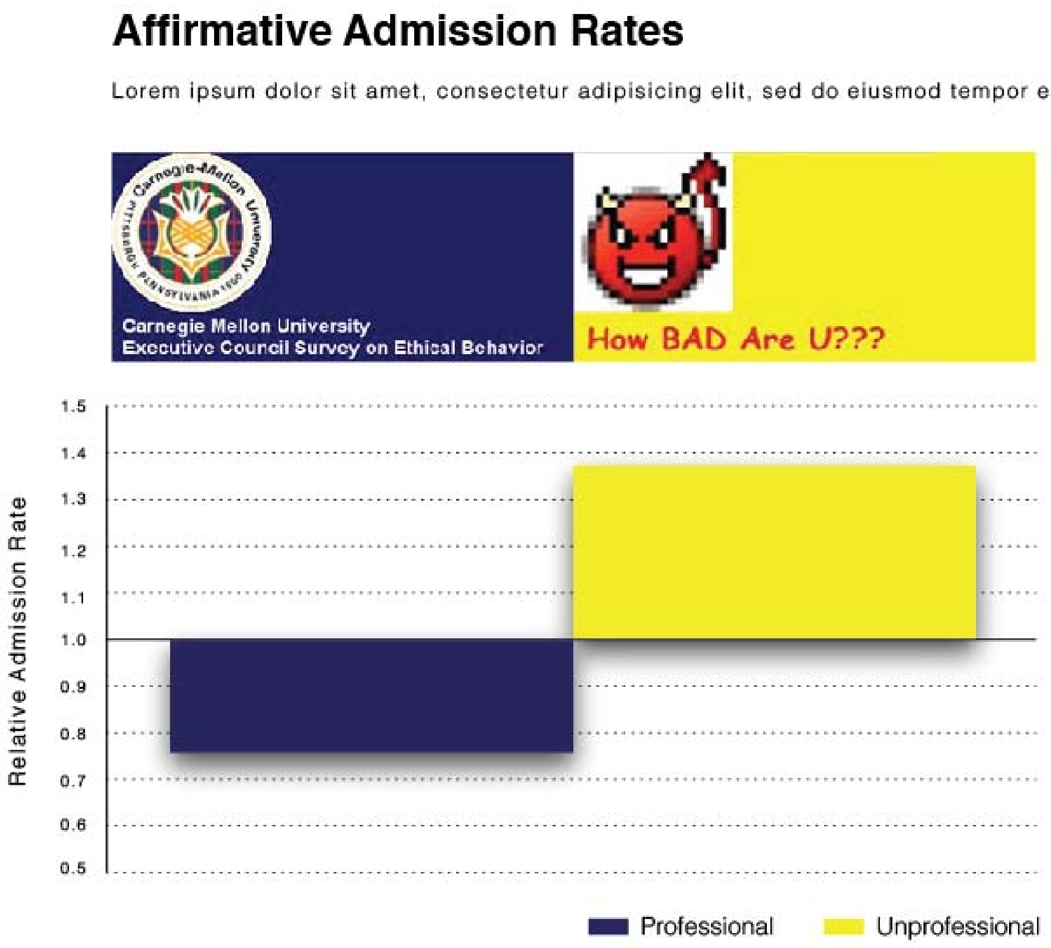

Figure 2. The impact of cues on disclosure behavior. A measure of privacy behavior often used in empirical studies is a subject's willingness to answer personal, sometimes sensitive questions – for instance, by admitting or denying having engaged in questionable behaviors. In an online experiment (*45*), individuals were asked a series of intrusive questions about their behaviors. Across conditions, the interface of the questionnaire was manipulated, to look more or less professional. The y axis captures the mean affirmative admission rates (AAR) to questions rates as intrusive (i.e., the proportion of questions answered affirmatively) normed, question by question, on the overall average AAR for the question. Subjects revealed more personal and even incriminating information on the website with a cheesier design, even though the site with the formal interface was judged by other respondents to be much safer, suggesting how cues can impact privacy behavior in manners that are unrelated, or even negatively related, to normative bases of decision making.

Changes in Facebook default profile settings over time (2005-2014). Degree of visibility of different fields of profiles based on default settings.

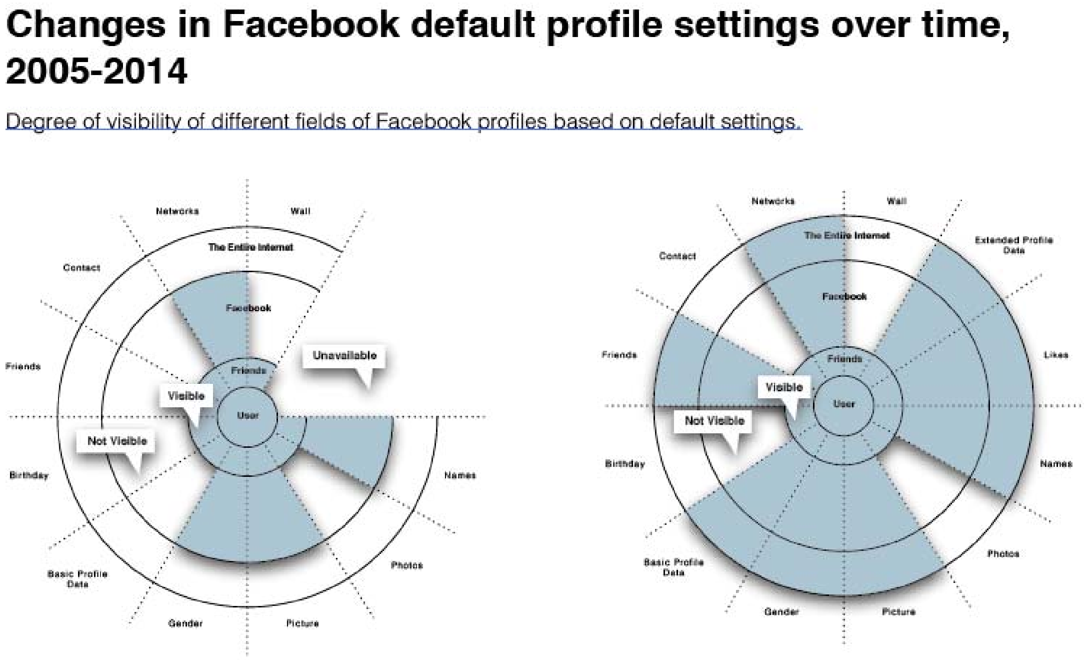

Changes in Facebook default profile settings over time, 2005-2014

Degree of visibility of different fields of Facebook profiles based on default settings.

Figure 3. Changes in Facebook default profile settings over time (2005-2014). Over time, Facebook profiles included an increasing amount of fields and, therefore, types of data. In addition, default visibility settings became more revelatory between 2005 and 2014, disclosing more personal information to larger audiences, unless the user manually overrode the defaults. This figure is based on the authors' data and the original visualization created by Matt McKeon, available at http://mattmckeon.com/facebook-privacy/.

Box 1. Privacy: A modern invention?

Is privacy a modern, burgeois, and uniquely Western invention? Or are privacy needs a universal feature of human societies? While *access to* privacy is certainly affected by socio-economic factors (Hargittai 2008) (some have referred to privacy as a 'luxury good' 27), and while privacy norms greatly differ across cultures (*63*, *84*), the *need for* privacy seems a common human trait. Scholars have uncovered evidence of privacy-seeking behaviors across peoples and cultures separated by time and space: from ancient Rome and Greece (Westin 1967, Aries et al 1992), to pre-industralized Javanese, Balinese, or Tuareg societies (Murphy 1954; Westin 1984). Privacy, as Altman (1977) noted, appears simultaneously culturally specific and culturally universal. Cues of a common human quest for privacy are also found in the texts of ancient religions: the Quran (49:12) instructs against spying on one another (Hayat 2007); the Talmud (Bava Batra 60a) advises to position windows so that they do not directly face those of one's neighbors (Enkin 2012); in the Bible (Genesis, 3:7), Adam and Eve discover their nakedness after eating the fruit of knowledge, and cover themselves in shame from the prying eyes of God (Rykwert 2001). (In addition, for a discussion of privacy in Confucian and Taoist cultures, see Whitman 1985.) Implicit in this heterogeneous selection of historical examples is the observation that there exist multiple notions of privacy. Notwithstanding their diversity, they all relate to the permeable yet profoundly consequential boundaries of public and private. While contemporary attention focuses on *informational* privacy, privacy has been also construed as territorial and physical, and linked to concepts as diverse as surveillance, exposure, intrusion, insecurity, appropriation, as well as secrecy, protection, anonymity, dignity, or even freedom (for a taxonomy, see Solove 2006). Such definitional diversity is reflected in the variety of literature streams we cover in this Review.

1. E. Hargittai, in *Social Stratification*, D. Grusky Ed. (Westview, Boulder, 2008) pp.936-113.
2. P. Ariès, P. Veyne, G. Duby, Eds., A History of Private Life: From Pagan Rome to Byzantium (Harvard Univ. Press, Cambridge, 1992)
3. R.F. Murphy, Social Distance and the Veil. *Am. Anthropol.* **66**, 1257-1274 (1964).
4. Westin, Alan. "The origins of modern claims to privacy." (1984), in Schoeman, Ferdinand David, ed. Philosophical dimensions of privacy: An anthology. Cambridge University Press, 1984.
5. Altman, Privacy regulation: Culturally universal or culturally specific?. *J. Soc. Iss.* **33**, 66-84 (1977).
6. M.A. Hayat, Privacy and Islam: From the Quran to data protection in Pakistan. *Information & Communications Tech. L.* **16**, 137-148 (2007).
7. Enkin, http://www.torahmusings.com/2012/07/privacy/
8. Rykwert, Joseph. "Privacy in antiquity." Social research (2001): 29-40.
9. Whitman, Christina B. "Privacy in Confucian and Taoist Thought." In Individualism and Holism: Studies in Confucian and Taoist Values, edited by D. Munro. Ann Arbor: Univ. of Michigan, Center for Chinese Studies, 1985.
10. D.J. Solove, A taxonomy of privacy. *Univ. Penn. L. Rev.* **154**, 477-564 (2006).