

# Privacy and human behavior in the age of information

Alessandro Acquisti,<sup>1\*</sup> Laura Brandimarte,<sup>1</sup> George Loewenstein<sup>2</sup>

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 are 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 e-mails, 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 in the cloud; navigate streets with geotracking systems; and 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. Accompanying the acceleration in data collection are steady advancements in the ability to aggregate, analyze, and draw sensitive inferences from individuals' data (2).

Both firms and individuals can benefit from the sharing of once hidden data and from the application of increasingly sophisticated analytics to larger and more interconnected databases (3). So too can society as a whole—for instance, when electronic medical records are combined to observe novel drug interactions (4). On the other hand, the potential for personal data to be abused—for economic and social discrimination, hidden influence and manipulation, coercion, or censorship—is alarming. The erosion of privacy can threaten our autonomy, not merely as consumers but as citizens (5). Sharing more personal data does not necessarily always translate into more progress, efficiency, or equality (6).

Because of the seismic nature of these developments, there has been considerable debate about individuals' ability to navigate a rapidly evolving privacy landscape, and about what, if anything, should be done about privacy at a policy level. Some trust people's ability to make self-interested

decisions about information disclosing and withholding. Those holding this view tend to see regulatory protection of privacy as interfering with the fundamentally benign trajectory of information technologies and the benefits such technologies may unlock (7). Others are concerned about the ability of individuals to manage privacy amid increasingly complex trade-offs. Traditional tools for privacy decision-making such as choice and consent, according to this perspective, 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 the social and behavioral sciences. We highlight factors that influence decisions to protect or surrender privacy and how, in turn, privacy protections or 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 socioeconomic power. Accordingly, this Review focuses on the concept of, and literature around, informational privacy (that is, privacy of personal data) but 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 (9).

We use three themes to organize and draw connections between streams of privacy research that, in many cases, have unfolded independently. The first theme is people's uncertainty about the nature of privacy trade-offs, and their own preferences over them. The second 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. Although 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 amplified by 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, or anonymity (10); by protecting personal information; but also through deception and dissimulation (11). 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 (12). Sometimes, these motivations are so visceral and primal that privacy-seeking behavior emerges swiftly and naturally. This is often the case when physical privacy is intruded—such as when a stranger encroaches in one's personal space (13–15) or demonstratively eavesdrops on a conversation. However, at other times (often including 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 or 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. First, whereas 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. Second, privacy is rarely an unalloyed good; it typically involves trade-offs (16). For example, ensuring the privacy of a consumer's

<sup>1</sup>H. John Heinz III College, Carnegie Mellon University, Pittsburgh, PA, USA. <sup>2</sup>Dietrich College, Social and Decision Sciences, Carnegie Mellon University, Pittsburgh, PA, USA.

\*Corresponding author. E-mail: [acquisti@andrew.cmu.edu](mailto:acquisti@andrew.cmu.edu)

purchases may protect her from price discrimination but also deny her the potential benefits of targeted offers and advertisements.

Elements that mitigate one or both of these exacerbating factors, by either increasing the tangibility of privacy harms or making trade-offs explicit and simple to understand, will generally affect privacy-related decisions. This is illustrated by one laboratory 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 (17). 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 (18) 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 (19). Although logically such assurances should lead to greater divulgence, they often had the opposite effect because they elevated respondents' privacy concerns, which without assurances would have remained dormant.

The remarkable uncertainty of privacy preferences comes into play in efforts to measure individual and group differences in preference for privacy (20). For example, Westin (21) 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 about privacy and express particular concern over losing control of their personal information or others gaining unauthorized access to it (22, 23). However, doubts about the power of attitudinal scales to predict actual privacy behavior arose early in the literature (24). 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 toward sharing data (25). Next, they 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 answering 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 (26, 27). Thus, one might care deeply about privacy in general but, depending on the costs and benefits prevailing in a specific situation, seek or not seek privacy protection (28).

This explanation for the privacy paradox, however, is not entirely satisfactory for two reasons. The first is that it fails to account for situations in which attitude-behavior dichotomies arise under high correspondence between expressed concerns and behavioral actions. For example, one study compared attitudinal survey answers to actual social media behavior (29). Even 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% revealed their political orientation, and 21% revealed 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 (16, 28); 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 United States, 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. Similar to 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 because technologies for information sharing have been enormously successful, whereas technologies for information protection have not, individuals hold overall low valuations of privacy. However, that is not always the case: Although individuals at times give up personal data for small benefits or discounts, at other times they 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. The cards could be used online or in stores, just like debit cards. Participants were given either 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). Initially, half of the participants were given one type of card, and half the other. Then, they were all offered the opportunity to switch. Some shoppers, for example, were given the anonymous \$10 card and were asked whether they would accept \$2 to "allow my name to be linked to transactions done with the card"; other subjects were asked whether they would accept a card with \$2 less value to "prevent my name from being linked to transactions done with the card." Of the subjects who originally held the less valuable but anonymous card, five times as many (52.1%) chose it and kept it over the other card than did those who originally held 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. Similar to 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 unimportance) are human motives no less fundamental than 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, increase self-esteem (36), and fulfill ego needs (37). In a series of functional magnetic resonance imaging 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. Indeed, subjects in one of the experiments were willing to forgo money in order to disclose about themselves (38).

### Context-dependence

Much evidence suggests that privacy is a universal human need (Box 1) (39). However, when people are uncertain about their preferences they often search for cues in their environment to provide guidance. And because 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, privacy fundamentalists, or privacy unconcerned, depending on time and place (40).

The way we construe and negotiate public and private spheres is context-dependent because the boundaries between the two are murky (41): 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, usually we may be more comfortable sharing secrets with friends, but at times we may reveal surprisingly personal information to a stranger on a plane (42). 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 (43). Thus, seeking privacy in public is not a contradiction; individuals can manage privacy even while sharing information, and even on social media (44). 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) (Fig. 1) (45).

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 the existence of some degree of privacy protection (46). In other situations, however, cues can be unrelated, or even negatively related,

to normative bases of decision-making. For example, in one online experiment (47) individuals were more likely to reveal personal and even incriminating information on a website with an unprofessional and casual 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 (Fig. 2). Yet in other situations, it is the physical environment that influences privacy concern and associated behavior (48), 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 rooms with bare cement and overhead fluorescent lighting (49).

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 (50). Observing other people reveal information increases the likelihood that one will reveal it oneself (51). 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, manipulated unbeknownst to them, about the percentage of other participants who in the same survey had admitted to having engaged in a given behavior. Being provided with information that 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 have found 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 (52). 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 (53) 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 (54), for instance, turbulence of privacy boundaries, as a result of 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, such as the introduction of a camera or other monitoring devices, privacy concern is likely to be activated. For instance, surveillance can produce discomfort (55) and negatively affect worker productivity (56). However, privacy concern, like other motivations, is adaptive; people get used to levels of intrusion that do not

### Fig. 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 publicly revealed their birthday in 2005, but less than 20% in 2011. The decreasing trend is not uniform, however. After decreasing for several years, the percentage of profiles that publicly revealed their high school roughly doubled between 2009 and 2010—after Facebook changed the default visibility settings for various fields on its profiles, including high school (bottom), but not birthday (top) (45).

### Disclosure behavior in online social media

Percentage of profiles publicly revealing information over time (2005-2011)



change over time. In an experiment conducted in Helsinki (57), the installation of sensing and monitoring technology in households led family members initially to change their behavior, 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 major implications for the risks associated with modern information and communication technology (58). 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 (59) where we become social media friends with our coworkers and post pictures to an indistinct flock of followers. With different social groups mixing on the Internet, separating online and offline identities and meeting our and others’ expectations regarding privacy becomes more difficult and consequential (60).

### Malleability and influence

Whereas 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 (61). 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.

Default settings are an important tool used by different entities to affect information disclosure. A large body of research has shown that default settings matter for decisions as important as organ donation and retirement saving (62). Sticking to default settings is convenient, and people often interpret default settings as implicit recommendations (63). Thus, it is not surprising that default settings for one’s profile’s visibility on social networks (64), or the existence of opt-in or opt-out privacy policies on websites (65), affect individuals’ privacy behavior (Fig. 3).

In addition to default settings, websites can also use design features that frustrate or even confuse users into disclosing personal information (66), a practice that has been referred to as “malicious interface design” (67). Another obvious strategy that commercial entities can use to avoid raising privacy concerns is not to “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 (68) and can respond with negative “reactance” (69).

Various so-called “antecedents” (70) 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 that are intended to protect privacy can establish trust, con-

cerns 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 (40), which suggests that simply posting a policy that consumers do not read may lead to misplaced feelings of being protected.

Control is another feature that can inculcate trust and produce paradoxical effects. Perhaps because of its lack of controversiality, control has

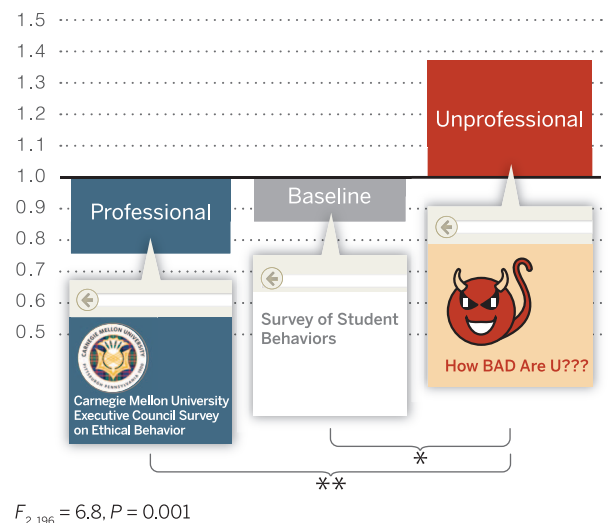
### Box 1. Privacy: A modern invention?

Is privacy a modern, bourgeois, and distinctly Western invention? Or are privacy needs a universal feature of human societies? Although access to privacy is certainly affected by socioeconomic factors (87) [some have referred to privacy as a “luxury good” (15)], and privacy norms greatly differ across cultures (65, 85), the need for privacy seems to be a universal human trait. Scholars have uncovered evidence of privacy-seeking behaviors across peoples and cultures separated by time and space: from ancient Rome and Greece (39, 88) to preindustrialized Javanese, Balinese, and Tuareg societies (89, 90). Privacy, as Altman (91) noted, appears to be 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 (92); the Talmud (Bava Batra 60a) advises home-builders to position windows so that they do not directly face those of one’s neighbors (93); the Bible (Genesis, 3:7) relates how Adam and Eve discovered their nakedness after eating the fruit of knowledge and covered themselves in shame from the prying eyes of God (94) [a discussion of privacy in Confucian and Taoist cultures is available in (95)]. Implicit in this heterogeneous selection of historical examples is the observation that there exist multiple notions of privacy. Although 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 [a taxonomy is provided in (9)].

**Fig. 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 (47), individuals were asked a series of intrusive questions about their behaviors, such as “Have you ever tried to peek at someone else’s e-mail without their knowing?” Across conditions, the interface of the questionnaire was manipulated to look more or less professional. The y axis captures the mean affirmative admission rates (AARs) to questions that were rated as intrusive (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 more casual design, even though the site with the formal interface was judged by other respondents to be much safer. The study illustrates how cues can influence privacy behavior in a fashion that is unrelated, or even negatively related, to normative bases of decision-making.

### A measure of privacy behavior


Relative admission rates in an experiment testing the impact of different survey interfaces on willingness to answer questions about various sensitive behaviors



**Fig. 3. Changes in Facebook default profile visibility 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 (top) and 2014 (bottom), disclosing more personal information to larger audiences, unless the user manually overrode the defaults (fields such as “Likes” and “Extended Profile Data” did not exist in 2005). “Basic profile data” includes hometown, current city, high school, school (status, concentration, secondary concentration), interested in, relationship, workplace, about you, and quotes. Examples of “Extended profile data” include life events such as new job, new school, engagement, expecting a baby, moved, bought a home, and so forth. “Picture” refers to the main profile image. “Photos” refers to the additional images that users might have shared in their account. “Names” refers to the real name, the username, and the user ID. This figure is based on the authors’ data and the original visualization created by M. McKeon, available at <http://mattmckeon.com/facebook-privacy>.

**Default visibility settings in social media over time**



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 (46), which in turn can have unintended effects. For instance, one study found that participants who were provided with greater explicit control over whether and how much of their personal information researchers could publish ended up sharing more sensitive information with a broader audience—the opposite of the ostensible purpose of providing such control (77).

Similar to the normative perspective on control, increasing the transparency of firms’ data practices would seem to be desirable. 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 (72), 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 (73). Indeed, and somewhat amusingly, it has been estimated that the aggregate opportunity cost if U.S. consumers actually read the privacy policies of the sites they visit would be \$781 billion/year (74).

Although uncertainty and context-dependence lead naturally to malleability and manipulation, not all malleability is necessarily sinister. Consider monitoring. Although monitoring can cause discomfort and reduce productivity, the feeling of being observed and accountable can induce people to engage in prosocial behaviors or (for better or for worse) adhere to social norms (75). Prosocial behavior can be heightened by monitoring cues as

simple as three dots in a stylized face configuration (76). By the same token, the depersonalization induced by computer-mediated interaction (77), either in the form of lack of identifiability or of visual anonymity (78), can have beneficial effects, such as increasing truthful responses to sensitive surveys (79, 80). Whether elevating or suppressing privacy concerns is socially beneficial critically depends, yet again, on context [a meta-analysis of the impact of de-identification on behavior is provided in (81)]. For example, perceptions of anonymity can alternatively lead to dishonest or prosocial behavior. Illusory anonymity induced by darkness caused participants in an experiment (82) to cheat in order to gain more money. This can be interpreted as a form of disinhibition effect (83), 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 (84).

**Conclusions**

Norms and behaviors regarding private and public realms greatly differ across cultures (85). Americans, for example, are reputed to be more open about sexual matters than are the Chinese, whereas the latter are more open about financial matters (such as 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. And 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 that 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 on to navigate the complex trade-offs 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, consumers, and citizens. Although such influence is not always or necessarily malevolent or dangerous, relinquishing control over one’s personal data and over one’s privacy 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 reviewed here suggest that policy approaches that rely exclusively on informing or “empowering” the individual are unlikely to provide adequate protection against the risks posed by recent information technologies. Consider transparency and control, two principles conceived as necessary conditions for privacy protection. The research we highlighted shows that they may provide insufficient protections and even backfire when used apart from other principles of privacy protection.

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—policies that include a baseline framework of protection, such as the principles embedded in the so-called fair information practices (86). People need assistance and even protection to aid in navigating what is otherwise a very uneven playing field. As highlighted by our discussion, a goal of public 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 real people—who are naïve, uncertain, and vulnerable—and should be sufficiently flexible to evolve with the emerging unpredictable complexities of the information age.

## REFERENCES AND NOTES

- V. Mayer-Schönberger, *Delete: The Virtue of Forgetting In the Digital Age* (Princeton Univ. Press, Princeton, 2011).
- L. Sweeney, *Int. J. Uncert. Fuzziness Knowl. Based Syst.* **10**, 557–570 (2002).
- A. McAfee, E. Brynjolfsson, *Harv. Bus. Rev.* **90**, 60–66, 68, 128 (2012).
- N. P. Tatonetti, P. P. Ye, R. Daneshjoui, R. B. Altman, *Sci. Transl. Med.* **4**, 125ra31 (2012).
- J. E. Cohen, *Stanford Law Rev.* **52**, 1373–1438 (2000).
- K. Crawford, K. Miltner, M. L. Gray, *Int. J. Commun.* **8**, 1663–1672 (2014).
- R. A. Posner, *Am. Econ. Rev.* **71**, 405–409 (1981).
- D. J. Solove, *Harv. Law Rev.* **126**, 1880–1903 (2013).
- D. J. Solove, *Univ. Penn. L. Rev.* **154**, 477–564 (2006).
- F. Schoeman, Ed., *Philosophical dimensions of privacy—An anthology* (Cambridge Univ. Press, New York, 1984).
- B. M. DePaulo, C. Wetzel, R. Weylin Sternglanz, M. J. W. Wilson, *J. Soc. Issues* **59**, 391–410 (2003).
- S. T. Margulis, *J. Soc. Issues* **59**, 243–261 (2003).
- E. Goffman, *Relations in Public: Microstudies of the Public Order* (Harper & Row, New York, 1971).
- E. Sundstrom, I. Altman, *Hum. Ecol.* **4**, 47–67 (1976).
- B. Schwartz, *Am. J. Sociol.* **73**, 741–752 (1968).
- R. S. Lauffer, M. Wolfe, *J. Soc. Issues* **33**, 22–42 (1977).
- J. Y. Tsai, S. Egelman, L. Cranor, A. Acquisti, *Inf. Syst. Res.* **22**, 254–268 (2011).
- P. Slovic, *Am. Psychol.* **50**, 364–371 (1995).
- E. Singer, H. Hippler, N. Schwarz, *Int. J. Public Opin. Res.* **4**, 256–268 (1992).
- V. P. Skotko, D. Langmeyer, *Sociometry* **40**, 178–182 (1977).
- A. Westin, Harris Louis & Associates, Harris-Equifax Consumer Privacy Survey (Tech. rep. 1991).
- M. J. Cullan, P. K. Armstrong, *Organ. Sci.* **10**, 104–115 (1999).
- H. J. Smith, S. J. Millberg, S. J. Burke, *Manage. Inf. Syst. Q.* **20**, 167–196 (1996).
- B. Lubin, R. L. Harrison, *Psychol. Rep.* **15**, 77–78 (1964).
- S. Spiekermann, J. Grossklags, B. Berendt, *E-Privacy in 2nd Generation E-Commerce: Privacy Preferences versus Actual Behavior* (Third ACM Conference on Electronic Commerce, Tampa, 2001), pp. 38–47.
- P. A. Norberg, D. R. Horne, D. A. Horne, *J. Consum. Aff.* **41**, 100–126 (2007).
- I. Ajzen, M. Fishbein, *Psychol. Bull.* **84**, 888–918 (1977).
- P. H. Klopfer, D. I. Rubenstein, *J. Soc. Issues* **33**, 52–65 (1977).
- A. Acquisti, R. Gross, in *Privacy Enhancing Technologies*, G. Danezis, P. Golle Eds. (Springer, New York, 2006), pp. 36–58.
- A. Acquisti, *Privacy in Electronic Commerce and the Economics of Immediate Gratification* (Fifth ACM Conference on Electronic Commerce, New York, 2004), pp. 21–29.
- I. Hann, K. Hui, S. T. Lee, I. P. L. Png, *J. Manage. Inf. Syst.* **24**, 13–42 (2007).
- A. Acquisti, L. K. John, G. Loewenstein, *J. Legal Stud.* **42**, 249–274 (2013).
- I. Altman, D. Taylor, *Social Penetration: The Development of Interpersonal Relationships* (Holt, Rinehart & Winston, New York, 1973).
- J. Frattaroli, *Psychol. Bull.* **132**, 823–865 (2006).
- J. W. Pennebaker, *Behav. Res. Ther.* **31**, 539–548 (1993).
- C. Steinfeld, N. B. Ellison, C. Lampe, *J. Appl. Dev. Psychol.* **29**, 434–445 (2008).
- C. L. Toma, J. T. Hancock, *Pers. Soc. Psychol. Bull.* **39**, 321–331 (2013).
- D. I. Tamir, J. P. Mitchell, *Proc. Natl. Acad. Sci. U.S.A.* **109**, 8038–8043 (2012).
- A. Westin, *Privacy and Freedom* (Athenäum, New York, 1967).
- C. J. Hoofnagle, J. M. Urban, *Wake Forest Law Rev.* **49**, 261–321 (2014).
- G. Marx, *Ethics Inf. Technol.* **3**, 157–169 (2001).
- J. W. Thibaut, H. H. Kelley, *The Social Psychology Of Groups* (Wiley, Oxford, 1959).
- H. Nissenbaum, *Privacy in Context: Technology, Policy, and the Integrity of Social Life* (Stanford Univ. Press, Redwood City, 2009).
- d. boyd, *It's Complicated: The Social Lives of Networked Teens* (Yale Univ. Press, New Haven, 2014).
- F. Stutzman, R. Gross, A. Acquisti, *J. Priv. Confidential.* **4**, 7–41 (2013).
- H. Xu, H. H. Teo, B. C. Tan, R. Agarwal, *J. Manage. Inf. Syst.* **26**, 135–174 (2009).
- L. K. John, A. Acquisti, G. Loewenstein, *J. Consum. Res.* **37**, 858–873 (2011).
- I. Altman, *The Environment and Social Behavior: Privacy, Personal Space, Territory, and Crowding* (Cole, Monterey, 1975).
- A. L. Chaikin, V. J. Derlega, S. J. Miller, *J. Couns. Psychol.* **23**, 479–481 (1976).
- V. J. Derlega, A. L. Chaikin, *J. Soc. Issues* **33**, 102–115 (1977).
- A. Acquisti, L. K. John, G. Loewenstein, *J. Mark. Res.* **49**, 160–174 (2012).
- Y. Moon, *J. Consum. Res.* **26**, 323–339 (2000).
- S. Petronio, *Boundaries of Privacy: Dialectics of Disclosure* (SUNY Press, Albany, 2002).
- 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.
- T. Honess, E. Charman, *Closed Circuit Television in Public Places: Its Acceptability and Perceived Effectiveness* (Police Research Group, London, 1992).
- M. Gagné, E. L. Deci, *J. Organ. Behav.* **26**, 331–362 (2005).
- A. Oulasvirta et al., *Long-Term Effects of Ubiquitous Surveillance in the Home* (ACM Conference on Ubiquitous Computing, Pittsburgh, 2012), pp. 41–50.
- 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.
- Z. Tufekci, *Bull. Sci. Technol. Soc.* **28**, 20–36 (2008).
- J. A. Bargh, K. Y. A. McKenna, G. M. Fitzsimons, *J. Soc. Issues* **58**, 33–48 (2002).
- R. Calo, *Geo. Wash. L. Rev.* **82**, 995–1304 (2014).
- E. J. Johnson, D. Goldstein, *Science* **302**, 1338–1339 (2003).
- C. R. McKenzie, M. J. Liersch, S. R. Finkelstein, *Psychol. Bull.* **117**, 414–420 (2006).
- 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.
- E. J. Johnson, S. Bellman, G. L. Lohse, *Mark. Lett.* **13**, 5–15 (2002).
- W. Hartzog, *Am. Univ. L. Rev.* **60**, 1635–1671 (2010).
- G. Conti, E. Sobiesk, *Malicious Interface Design: Exploiting the User* (19th International Conference on World Wide Web, ACM, Raleigh, 2010), pp. 271–280.
- A. Goldfarb, C. Tucker, *Mark. Sci.* **30**, 389–404 (2011).
- T. B. White, D. L. Zahay, H. Thorbjørnsen, S. Shavitt, *Mark. Lett.* **19**, 39–50 (2008).
- H. J. Smith, T. Dinev, H. Xu, *Manage. Inf. Syst. Q.* **35**, 989–1016 (2011).
- L. Brandimarte, A. Acquisti, G. Loewenstein, *Soc. Psychol. Personal. Sci.* **4**, 340–347 (2013).
- C. Jensen, C. Potts, C. Jensen, *Int. J. Hum. Comput. Stud.* **63**, 203–227 (2005).
- 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.
- A. M. McDonald, L. F. Cranor, *I/S: J. L. Policy Inf. Society.* **4**, 540–565 (2008).
- C. Wedekind, M. Milinski, *Science* **288**, 850–852 (2000).
- M. Rigdon, K. Ishii, M. Watabe, S. Kitayama, *J. Econ. Psychol.* **30**, 358–367 (2009).
- S. Kiesler, J. Siegel, T. W. McGuire, *Am. Psychol.* **39**, 1123–1134 (1984).
- A. N. Joinson, *Eur. J. Soc. Psychol.* **31**, 177–192 (2001).
- S. Weisband, S. Kiesler, *Self Disclosure On Computer Forms: Meta-Analysis And Implications* (SIGCHI Conference on Human Factors in Computing Systems, ACM, Vancouver, 1996), pp. 3–10.
- R. Tourangeau, T. Yan, *Psychol. Bull.* **133**, 859–883 (2007).
- T. Postmes, R. Spears, *Psychol. Bull.* **123**, 238–259 (1998).
- C. B. Zhong, V. K. Bohns, F. Gino, *Psychol. Sci.* **21**, 311–314 (2010).
- J. Suler, *Cyberpsychol. Behav.* **7**, 321–326 (2004).
- A. F. Shariff, A. Norenzayan, *Psychol. Sci.* **18**, 803–809 (2007).
- B. Moore, *Privacy: Studies in Social and Cultural History* (Armonk, New York, 1984).
- Records, Computers and the Rights of Citizens* (Secretary's Advisory Committee, U.S. Dept. of Health, Education and Welfare, Washington, DC, 1973).
- E. Hargittai, in *Social Stratification*, D. Grusky Ed. (Westview, Boulder, 2008), pp. 936–113.
- P. Ariès, G. Duby (Eds.), *A History of Private Life: From Pagan Rome to Byzantium* (Harvard Univ. Press, Cambridge, 1992).
- R. F. Murphy, *Am. Anthropol.* **66**, 1257–1274 (1964).
- A. Westin, in *Philosophical Dimensions of Privacy: An Anthology*, F. D. Schoeman Ed. (Cambridge Univ. Press, Cambridge, UK, 1984), pp. 56–74.
- I. Altman, *J. Soc. Issues* **33**, 66–84 (1977).
- M. A. Hayat, *Inf. Comm. Tech. L.* **16**, 137–148 (2007).
- A. Enkin, “Privacy,” [www.torahmusings.com/2012/07/privacy](http://www.torahmusings.com/2012/07/privacy) (2014).
- J. Rykwert, *Soc. Res. (New York)* **68**, 29–40 (2001).
- C. B. Whitman, in *Individualism and Holism: Studies in Confucian and Taoist Values*, D. J. Munro, Ed. (Center for Chinese Studies, Univ. Michigan, Ann Arbor, 1985), pp. 85–100.

## ACKNOWLEDGMENTS

We are deeply grateful to the following individuals: R. Gross and F. Stutzman for data analysis; V. Marotta, V. Radhakrishnan, and S. Samat for research; W. Harsch for graphic design; and A. Adams, I. Adjerid, R. Anderson, E. Barr, C. Bennett, R. Boehme, R. Calo, J. Camp, F. Cate, J. Cohen, D. Cole, M. Cullan, R. De Wolf, J. Donath, S. Egelman, N. Ellison, S. Fienberg, A. Forget, U. Gasser, B. Gellman, J. Graves, J. Grimmelmann, J. Grossklags, S. Guerses, J. Hancock, E. Hargittai, W. Hartzog, J. Hong, C. Hoofnagle, J. P. Hubaux, K. L. Hui, A. Joinson, S. Kiesler, J. King, B. Knijnenburg, A. Kobsa, B. Kraut, P. Leon, M. Madden, I. Meeker, D. Mulligan, C. Olivola, E. Peer, S. Petronio, S. Preibusch, J. Reidenberg, S. Romanosky, M. Rotenberg, I. Rubenstein, N. Sadeh, A. Sasse, F. Schaub, P. Shah, R. E. Smith, S. Spiekermann, J. Staddon, L. Strahilevitz, P. Swire, O. Tene, E. VanEpps, J. Vitak, R. Wash, A. Woodruff, H. Xu, and E. Zeide for enormously valuable comments and suggestions.

10.1126/science.aaa1465