

Secrets and Likes: The Drive for Privacy and the Difficulty of Achieving It in the Digital Age

Alessandro Acquisti, Laura Brandimarte, and George Loewenstein¹

Authors' version. Not final

Abstract

We review different streams of social science literature on privacy with the goal of understanding consumer privacy decision making and deriving implications for policy. We focus on psychological and economic factors influencing both consumers' desire and consumers' ability to protect their privacy, either through individual action or through the implementation of regulations applying to firms. Contrary to depictions of online sharing behaviors as careless, we show how consumers fundamentally care about online privacy, and present evidence of numerous actions they take to protect it. However, we also document how prohibitively difficult it is to attain desired, or even desirable, levels of privacy through individual action alone. The remaining instrument for privacy protection is policy intervention. However, again for both psychological and economic reasons, the collective impetus for adequate intervention is often countervailed by powerful interests that oppose it.

¹ Acknowledgments: We thank Aradhna Krishna and two anonymous reviewers for their sharp and insightful guidance. We thank Ross Anderson, Julie Cohen, Jessica Colnago, Lorrie Cranor, Daphne Chang, Judith Donath, Alisa Frik, Jim Graves, Chris Hoofnagle, Wolfgang Kerber, Steve Margulis, Kirsten Martin, Eyal Peer, Sandra Petronio, Elissa Redmiles, Florian Schaub, Daniel Solove, and Daniel Smullen for their time and their helpful comments. Some arguments in Section 4 are based on remarks prepared by one of the authors for the Atlanta FED 2019 Conference on Markets and the Economy. Acquisti gratefully acknowledges support via a grant from the Alfred P. Sloan Foundation.

1. Introduction

In the summer of 2016, a Facebook app featuring a personality test funneled millions of consumers' personal data to political operatives for use in targeted ads intended to influence voting in the upcoming US presidential elections. When this scheme was revealed, outrage ensued (Rosenberg et al., 2018). Many users who had filled out the test had not realized whom their data could be shared with or for what purpose (Chen, 2018). Yet, the threatened mass exodus from, and regulation of, social media never came (Wong, 2019). Four years later, little seems to have changed.

The Cambridge Analytica scandal encapsulates the modern privacy problem: seemingly carefree online sharing behaviors; concerns with misuses of data, quickly forgotten; and staggering—and yet so hard to demonstrate or quantify—individual and societal consequences. To make sense of this problem, we review diverse streams of social science literature and take account of both the psychological characteristics of individuals as well as the economic environment in which they attempt to navigate privacy issues.

An often-repeated claim is that, under the assault of novel information technologies, privacy is dead—both in the sense that it is no longer achievable and that concern about privacy is fading. In Section 2, we provide evidence against the latter assertion, including survey responses, field observations, and experimental results, all showing that consumers fundamentally care about privacy and often act on that concern. Actions to protect privacy are in fact so ubiquitous and second nature that they often go unnoticed.

If people do care about privacy and take actions to protect it, can they do so effectively in an increasingly digitized world? And, should they? In Sections 3 and 4, we discuss how, while the desire for privacy may be widespread, consumers' ability to reach desired or even desirable levels of privacy is far less so.

Section 3 focuses on psychological and behavioral factors limiting privacy attainment. Although consumers display a concern about privacy in a wide array of behaviors, there are

broad realms of life in which such concern is difficult to discern. Some motivations—such as the powerful desire to share information with other people—counteract the coexisting desire for protection. And some other features of human psychology—such as the tendency to adapt to a gradually evolving situation, glean social norms from other people’s behavior, and take greater risks when one feels in control—can reduce or even neutralize privacy concerns.

Section 4 focuses on economic factors that explain the difficulty, or even inability, of consumers to satisfy their desire for privacy—or, when such desire is absent but would be justified, in achieving beneficial levels of privacy. We also consider how the economic side of the equation overlaps with the psychological side—that is, the different ways that Internet-age firms exploit consumers’ psychological propensities to disclose. Furthermore, while some have proposed that the benefits that have accompanied the loss of privacy ultimately justify the costs, we examine whether the untroubled acceptance of the erosion of privacy may ignore real costs and consequences that threaten not only individuals, but also society. We conclude that market-based approaches to privacy, such as those in the United States, that rely on “notice and consent” strategies and consumer “responsibilization,” have failed to secure either desired, or desirable, levels of privacy, and that a combination of psychological and economic hurdles explains otherwise seemingly paradoxical patterns of privacy choice (Section 5).

Given that people should and do care about privacy, but cannot adequately manage it, what should be done? The inescapable conclusion, we argue in Section 6, is that the restoration of privacy can only be a matter of public policy. And yet, while numerous regulatory efforts have taken place around the world, economic and psychological factors again explain why adequate levels of protection remain hard to attain.

While these conclusions may appear pessimistic, we end by highlighting how the resiliency—and apparent universality—of a human drive for privacy provides hope for a future that balances privacy with sharing and data utility. We offer suggestions on research directions to investigate that balance.

2. Do Consumers Care About Privacy?

Do consumers care about, and act to protect, their privacy? Among Americans, the evidence for elevated privacy *concerns* has been ample and enduring. It can be found in Alan Westin's seminal surveys from the last century, which identified an overwhelming majority of Americans as troubled by potential misuses of their information (Westin, 2001), up to the most recent polls, showing a majority of Americans concerned about data harvesting by corporations (Pew, 2019). Yet, a common response to the above question is that, although people say that they care, they do not actually do—as proven by seemingly careless online disclosures (Johnson, 2020; Miller, 2014). Privacy is no longer a social norm, the narrative goes, or, in fact, privacy is dead (Sprenger, 1999).

We disagree. We show in this section that consumers do not just say they care about privacy, but in fact often take action to protect it. Surveys, field studies, and experiments—as well as common sense—show that consumers engage in privacy-regulating behaviors continually and in both online and offline scenarios, crossing the many diverse dimensions and definitions of privacy (Solove, 2005), from spatial to informational. The drive for privacy is underappreciated in part because individual actions to protect privacy are so ubiquitous and second nature that they go unnoticed, or are not construed as privacy behaviors.

2.1 The Evidence for Privacy-Seeking Behavior

In our everyday lives in the offline world, we instinctively and continually engage in privacy behaviors without even thinking: lowering our voice during intimate conversations; leaving a group of people to take a personal call; tilting a document we are reading so it is protected from prying eyes; and drawing curtains to ensure privacy in our bedrooms. Altman (1975) argued that privacy behaviors are so ubiquitous and common that they occur with little conscious awareness. Writing decades before the Internet age—and so, focusing on personal space rather than information—he noted that protection of personal space is instinctive and universal, across time and geography: Transgressions of personal space invoke a variety of

reactions, including a shifting of gaze, breaking eye contact, turning the body, or adopting protective postures (Altman, 1977).

Altman's insights apply also to how people interact on the Internet (Palen and Dourish 2003). Online, too, we engage continuously in behaviors that delimit the contours of our closeness or openness to others. Multiple times per day, we alternate between different email accounts or online personae to separate personal from professional spheres; pick privacy settings to manage the visibility of our social media posts; reply privately to group messages, carefully selecting recipients for our responses; enter (or rely on previously stored) passwords to keep information in our online accounts private; set "I am busy!" notices on instant messaging profiles to tell people not to contact us, right now; and turn on and off camera or audio on conference calls, depending on what we want to (or must) show to others. The precise motivations behind these behaviors are diverse; their common trait is the individual's attempt to regulate the boundaries of their interactions with others.

Observations of online privacy-seeking behaviors go well beyond the anecdotal. In the next subsections, we catalogue nonexhaustive examples of research evidence.

2.1.1 Surveys

Consider a seminal survey by Pew (2013): It found that an overwhelming majority (86%) of surveyed US adults reported having taken steps online to remove or mask their digital footprints. The steps were diverse, from less to more sophisticated, including clearing cookies, encrypting email, avoiding using their name, and using virtual private networks.

Numerous other surveys and interview reports provide similar evidence. For instance, a 2012 survey (Pew, 2012) found that the majority (58%) of social network site users had restricted access to their profiles. Another Pew (2015) survey found that among respondents who had heard of government surveillance, 34% had changed the way they communicated online, including avoiding using certain terms or uninstalling apps from their devices. In 2019, a survey of over 2,600 adult respondents in 12 of the world's largest economies conducted by

CISCO found that one third of them had switched companies or providers over their data policies or data sharing practices (CISCO 2019). A survey of 1,500 Americans by Instamotors found, in 2017, that 89% of respondents had taken at least one step to protect data—including changing passwords, not displaying personal information on social media, using password generators, covering laptop cameras when not in use, or using VPNs when browsing (Lewis, 2017). And a survey of 5,710 US participants conducted by DuckDuckGo in 2017 found that 46% had used “private browsing” (a browser setting to remove certain traces of browsing activities from a computer) at least once. In an in-depth interview study, Kang et al. (2015) found that 77% of nontechnical participants reported taking some action to hide or delete their digital traces, including using anonymous search engines and blockers for cookies and other trackers. And recent work on reactions to contact tracing suggests that 53% of the US population may be unwilling to install an app that would provide information to public health officials to track the spread of SARS-CoV-2 (KFF, 2020).

2.1.2 Field Studies

Self-reports do not always reflect actual behavior. Yet, observational field studies of consumer choice also provide evidence of privacy-seeking behavior. The behaviors encompass the many diverse scenarios in which privacy (and privacy invasions) plays a role in consumers’ lives: telemarketing annoyance, government surveillance, online tracking, social media intrusions, and so on.

For instance, by 2007, after the Federal Trade Commission had opened the National Do Not Call Registry (a database with the telephone numbers of individuals who do not want telemarketers to contact them), 72% of Americans had registered on the list (Bush, 2009). And following the revelations Edward Snowden made in 2013 regarding secret government surveillance programs, US consumers became less likely to read Wikipedia articles perceived as privacy sensitive (Penney, 2016).

Consumers try to elude online tracking in many different ways. Forty-one percent of 451 participants in an academic study were found to have used a feature called “private browsing” (which allows hiding some information about one’s browsing behaviors) at least once (Habib et al., 2018). A recent study of a group of consumers’ consent to firm targeting following the enactment of the European General Data Protection Regulation (GDPR) found that opt-in consent to the collection of intrusive data such as location information was as low as 5.5% (Godinho de Matos & Adjerid, 2019).

Social media users employ a variety of other strategies to carve out private spaces in online public fora, including self-censoring (Vitak & Ellison, 2013), “social steganography” (the act of hiding information in plain sight, by encoding messages for friends in otherwise public channels in ways that will allow only the intended recipients to discern their true meaning; Boyd & Marwick, 2011), and selective sharing (Vitak & Kim, 2014). For instance, Facebook users who were members of the Carnegie Mellon University’s network in 2005 progressively transitioned toward less public sharing of personal information over time: While 86% of those users were publicly sharing their date of birth in 2005, the percentage of them doing so decreased year by year, down to 22% in 2009 (Stutzman et al., 2013). Additionally, while only a miniscule proportion of Facebook users on that same Facebook network had altered their (highly visible) default search and visibility settings in 2005 (Gross & Acquisti, 2005), just about 7% of Facebook users studied by Fiesler et al. (2017) had *not* changed their default privacy settings by 2017. While the two samples are different, and self-selection cannot be excluded, the disparity in choices over time is stark.

2.1.3 Experiments

Evidence of privacy-seeking behavior arises from laboratory and field experiments as well. Many of them focus on exchanges involving privacy of data and monetary rewards. For instance, a field experiment testing for a gap in willingness to pay/willingness to accept for privacy showed that over 50% of participants were *not* willing to exchange a \$10 anonymous gift

card for a \$12 trackable one—essentially refusing a 20% increase in rewards to give away information on their purchasing decisions (Acquisti, John, & Loewenstein, 2013). When information about sellers' privacy policies was made salient and understandable, subjects in a laboratory were willing to pay roughly a 4% premium to purchase from more privacy-oriented sellers (Tsai et al., 2011). In a survey-based experiment, making privacy policies salient reduced subjects' disclosure of personal information, regardless of whether those policies were protective or intrusive of one's privacy (Marreiros et al., 2017). Online shoppers in a field experiment were willing to pay approximately one euro to keep their mobile phone number private (Jentzsch et al., 2012). In an online experiment in which subjects were first asked to answer survey questions in exchange for monetary rewards, and later asked to share their Facebook profiles with the researchers for an additional bonus, a majority were not willing to share their data for \$2.50 (Svirsky, 2019). In a choice experiment, subjects were willing to make one-time payments ranging from around \$1 to over \$4 in order to prevent each smartphone app from accessing information such as browsing history, contacts, location, or texts (Savage & Waldman, 2015).

2.1.4 A Note About Samples

Most of the examples above are based on WEIRD (Western, Educated, Industrialized, Rich, and Democratic) samples (Henrich et al., 2010). However, the evidence of “privacy”-seeking behaviors across cultures and historical periods is ample (Moore, 1984; Murphy, 1964), from the use of secret paths in the woods by some tribes in Brazil, to Javanese people hiding their emotions and speaking softly, and from the rearranging of the huts by the Pygmies of Zaire in response to the arrival of new people to the camp, to the covering of almost their entire face by the Tuareg of Northern Africa (Altman, 1977). As Altman concluded, privacy is simultaneously culturally specific and culturally universal.

2.2 Acknowledging Complexities, and Reconciling the Evidence

Although we have deliberately highlighted evidence of privacy-seeking behaviors to counter a prevailing narrative of disappearing concerns, there exists ample evidence of people *not* bothering to protect information, or engaging publicly in behaviors only a short time ago considered highly private. There is evidence of consumers being unwilling to pay for data protection (Beresford et al., 2012; Preibusch et al., 2013), and choosing to give up personal data in exchange for small convenience and small rewards (Athey et al., 2017). Although by now it is well known that mobile apps collect and share sensitive information with third parties (Almuhimedi et al., 2015), the number of app downloads increases every year (Statista, 2016). Major data breaches have become common (Fiegerman, 2017), yet most people seem willing to trust companies with personal information in exchange for relatively small benefits (Ghose, 2017). And a plethora of widespread, readily observable, everyday online behaviors seem to bespeak an overall lack of concern.

Evidence of disclosure-seeking behavior on its own, however, does not contradict the argument that consumers care for privacy. First, the work of Altman (1975, 1977) serves as an antidote to simple notions of privacy as a static condition of withdrawal, protection, or concealment of information (e.g., Posner, 1978). Altman construed privacy as a dialectic and dynamic process of *boundary regulation*. Privacy regulation encompasses *both* the opening *and* the closing of the self to others. By balancing and alternating the two, individuals manage interpersonal boundaries to achieve desired levels of privacy—an optimal amount of social interaction, or an optimal balance of both disclosing and protecting personal information. Privacy regulation is thus dynamic—a process highly dependent on and responsive to changes in context and a process that applies equally to the many diverse dimensions of privacy the literature has explored (a diversity this manuscript embraces, as evidenced by the disparate consumer scenarios it covers). Consistent with this account (see also Petronio 2002), the seemingly contrasting examples of protection- and disclosure-seeking behaviors illustrate how,

while we *manage* privacy *all the time*, we do not continuously *protect* our data. It would be undesirable (and unfeasible) for any individual to do so.

Second, the evidence for seemingly privacy-neglecting behaviors highlights a deeper issue: Privacy is extraordinarily difficult to manage, or regulate, in the Internet age. Consider, again, some of the examples of protective behaviors presented earlier in this section; they have a second side. Although Facebook users on the Carnegies Mellon University's network *did* become less likely to share their personal information with strangers between 2005 and 2009, changes in the Facebook interface in late 2009 and early 2010, affecting the visibility of certain profile fields, abruptly reversed that protective trend, making *public* again, for a significant portion of users, personal information those users had attempted to keep private (Stutzman et al., 2013). Likewise, although a DuckDuckGo survey *did* find that a substantial proportion of Internet users had tried to use private browsing, it also found that two-thirds of those users misunderstood (in fact, overestimated) the degree of protection that private browsing provided. And, while the cited Instamotor survey *did* find that 89% of respondents had taken at least one step to protect data, the percentage of respondents taking *all* of those steps was exceedingly small; in fact, some of the more protective steps (such as using VPNs) were adopted by a small minority. Not coincidentally, those steps were also the ones less familiar to average users, and costlier to adopt.

This more nuanced look at the evidence suggests that claims concerning privacy being “dead” too often confuse *wants* with *opportunities*—what people want and what they can actually achieve. The desire for privacy may well be virtually universal; consumers, offline as well as online, continually attempt to regulate boundaries between public and private. Yet, the opportunities and ability to do so effectively—to achieve desired levels of privacy—may be shrinking. As the above examples suggest, and as the next sections demonstrate, the reasons are both psychological (Section 3) and economic (Section 4).

3: Can Individuals Effectively Manage Privacy Online?

Economists use the term “revealed preferences” to refer to how consumers’ true valuations can be revealed by their behavior, such as their choices in the marketplace. Applied to the privacy domain, a revealed preference argument would posit that consumers protect and share precisely what they desire—from disclosing personal information on social media to covering online footprints using privacy-enhancing technologies. The choices they make, according to this perspective, should be optimal for them personally and for society as a whole: If privacy behaviors express true preference for privacy, market outcomes based on consumers’ choices would, in the absence of externalities, maximize aggregate utility.

On the surface, such reasoning appears consistent with the Altmanian notion of privacy as a process of boundary regulation, according to which the individual deliberately chooses when and what to protect or to share. In reality, regulating one’s privacy is aspirational: In Altman’s terms, *desired* privacy may not be matched by *achieved* privacy, and market behaviors may not always necessarily capture underlying preferences for privacy. We focus, in this section, on some psychological factors causing the discrepancy.

3.1 Consumer Characteristics Affecting Privacy Behavior

The privacy literature has increasingly drawn from research in psychology and behavioral economics to provide empirical evidence of numerous processes affecting, and sometimes impairing, privacy-related decision making (Margulis 2003). These factors range from privacy “calculus” to emotions; from asymmetric information to bounded rationality; and from resignation and learned helplessness to cognitive and behavioral biases (Acquisti, Brandimarte, & Loewenstein, 2015). Some of them are summarized in Table 1 and discussed in the rest of this section. Together, they explain when privacy-related behaviors capture actual preferences, and when they may not.

Table 1: Some psychological factors affecting privacy decision making

Psychological factor	Description	Representative consequence	Firms' response
Information asymmetries	Consumers are unaware of the diverse ways firms collect and use their data	Consumers cannot respond to risks they are unaware of	Increases firms' ability to collect and use consumer information
Bounded rationality	Consumers lack the processing capacity to make sense of the complexities of the information environment	Few read, or even could make sense of, privacy policies	Writing policies using sophisticated, legalistic terms that obscure the central issues
Present bias	Overemphasizing immediate, and under-weighting delayed, costs and benefits	Consumers will incur long-term costs—for example, intrusive profiling and advertising—in exchange for small immediate benefits—for example, online discounts	Offering small benefits in exchange for consumer data sharing
Intangibility	Putting little weight on outcomes that are intangible—difficult to isolate or quantify	Consequences of privacy violations are often diffuse and difficult to connect with specific actions	Making it difficult for consumers to draw connections between specific acts of data sharing and specific privacy violations (e.g., price discrimination)
Constructed preferences	Uncertainty about one's preferences leads people to rely on crude decision heuristics that often run counter to considerations of objective costs and benefits	Sticking with default privacy settings	Setting defaults that are advantageous to the firm rather than to the consumer
Illusory control	The feeling (often illusory) that one is in control of a situation leads to increased risk-taking	Consumers share more when given more granular control over privacy settings	Provide consumers with more granular privacy controls to encourage disclosure
Herding	The tendency to imitate the behavior of other people	Consumers share more when they see others sharing more on social media	Provide social media feeds that convey a maximal sense of others' sharing
Adaptation	The tendency to get used to risks that are unchanged over time or that increase gradually	Despite ever-increasing violations of privacy, consumers adapt to them and accept them	Change data usage practices gradually

The drive to share	The powerful drive to share information, including personal information	Sharing of highly private, or even incriminating, information (e.g., on social media)	Working behavioral levers that elicit the motive to share (e.g., recommending photographs to share)
--------------------	---	---	---

3.1.1 Informational Asymmetries and Bounded Rationality

Perhaps the most obvious reason why consumers cannot achieve desired levels of privacy is ignorance about how their information is collected, disseminated, and used. What one does not know *can* hurt one, and worse, negative outcomes one is unaware of are impossible to avoid. Most Americans are unaware of the data practices of firms, and very few are capable of making sense of the ubiquitous privacy policies they blithely agree to in order to access different apps and websites (Pew, 2019). Even those who do have the expertise to make sense of such policies are unlikely to have the time to read through them. One amusing study estimated the annual national opportunity cost of Americans actually reading privacy policies at over three quarters of a trillion dollars (McDonald & Cranor, 2008). And many Americans believe that privacy policies provide them with protections, when the reverse is more likely to be true; they provide firms with uninformed consent to use and often sell their information (Hoofnagle & Urban, 2014). Moreover, people may derive (incorrect) assumptions of protection based on expected norms or privacy policies (Martin & Nissenbaum, 2016). For instance, even privacy-savvy consumers who use tools such as private browsing overestimate the degree of protection they obtain (Habib et al., 2018).

3.1.2 Present Bias and Intangibility

A second reason why people may take actions that do not realize their desired levels of privacy is that the benefits of actions such as using an app or accessing a source of information are typically immediate, while the privacy costs are often delayed, uncertain, and intangible (Acquisti, 2004). Work in behavioral economics shows that consumers respond

disproportionately to costs and benefits that are immediate, a phenomenon known as “present-bias.” As a result, minor inconveniences, such as the time and effort of completing a form, can have a huge impact on behavior (Bettinger et al., 2012). People are also very bad at dealing with small probabilities of negative outcomes (see, e.g., Kunreuther et al., 1978). In some cases (e.g., the threat of terrorism), they overweigh such small probabilities, but in many other situations (e.g., when it comes to insuring against earthquakes) they simply ignore them. Privacy is virtually a “perfect storm” when it comes to all these considerations. The benefits of engaging in privacy-challenging activities (e.g., online discounts offered in exchange of personal information) are typically immediate and tangible, while the consequences are typically delayed, and have either low likelihood (e.g., catastrophic identity theft, or loss of job due to social media postings made while in college), or have very high likelihood but are small and intangible (delays in the loading time of a website due to trackers and ads).

3.1.3 Constructed Preferences

Other research in psychology and behavioral economics shows that when consumers are uncertain about their own preferences, they seize upon any available cues to aid them in making a decision, a process known as “constructed preferences” (Slovic, 1995). In one study documenting the phenomenon of constructed preferences for privacy, John et al. (2011) asked students to answer sensitive questions (e.g., if they had ever peeked at someone else’s email without them knowing), or even potentially incriminating questions (e.g., if they had driven while under the influence), using either an official-looking university interface which provided assurances of anonymity and confidentiality, or an obviously amateur interface. Although an independent group of raters evaluated the professional website as much more secure, students were considerably more likely to admit to these behaviors on the unprofessional website. Cues that most evaluators agreed signaled safety of divulging, in fact led respondents to “clam up”—presumably because the formal website reminded them of the potential hazards of excessive divulgence. The features of the device utilized to self-disclose also matter. People seem more

comfortable with, and end up disclosing more intimate thoughts on social media, when using smartphones than PCs (Melumad & Meyer, 2020). Mental accounting also plays an important role when deciding whether to share personal information with marketers (White et al., 2014).

3.1.4 Illusory Control

In another line of research, the three of us (Brandimarte et al., 2013) documented a “control paradox.” Much as people feel safer driving than flying because they feel more in control when they drive, the illusion that they have greater control over their privacy leads consumers to divulge more freely, even though they rarely actually choose to exert that control. In one study from that paper, students were asked sensitive questions, such as whether they had ever stolen anything, and were randomly assigned to four conditions that varied the degree of control they perceived over the publication of their responses. In one condition, they were told that if they answered all their answers would be published on a university website. In a second, offering a modicum of greater control, they were asked to check a box to allow publication of all their answers. A third condition that increased control even further asked them to check a box for each question to allow publication of that answer. And a fourth condition was the same as the second (in which a single box led to publication of all answers), but respondents were also asked to provide demographic information that would have allowed us to uncover the identity of the respondent. The striking results from the study were that, although largely illusory, greater control led to greater admission of sensitive or illegal behaviors, but collecting detailed demographic information, which should have given rise to fears of reidentification, had no impact.

The flip side of feeling in control is feeling that one has no control over a situation, and perversely, this can also lead to a loss of vigilance about privacy issues. Draper and Turow (2019) propose that people’s apparent inaction regarding privacy issues is the result of a sense of helplessness they refer to as “digital resignation.” As much as people wish to manage the river of personal information that flows from them to companies, they realize that it is outside of

their capabilities, so they simply give up. Barassi (2019) notices that being an active citizen in the modern world necessarily implies digital participation (even by children, whose parents post photographs, videos, and all kinds of other information, leaving children no choice in the matter), which inevitably leads to one's "datafication."

3.1.5 Herding

One of the most common cues that people use to decide how much information to reveal is what others reveal. Documenting this and a variety of related effects, Acquisti, John, and Loewenstein (2012) conducted a study in which a panel of New York Times readers were asked a series of sensitive questions and, after answering each, were given the ostensible distribution of answers by other participants. There were three possible answers: They had engaged in the behavior, they had not engaged in the behavior, or they chose not to reveal whether they had. In each of the three conditions, the feedback they got about other people's responses was manipulated to make it look as if most people gave one of these three responses. Confirming the influence of other people, over time individuals' responses trended toward echoing those of other people. When other people denied engaging in the behaviors, or reported unwillingness to answer, then respondents began, themselves, to gravitate toward greater concealment; but when others admitted to the behaviors, respondents were more willing to do so themselves.

3.1.6 Adaptation

Consumers' responses to problems, including violations of privacy, have another pernicious property: They are adaptive. Problems that seem acute, and initially draw a lot of attention and alarm, tend to recede into the background, even if they remain constant over time or worsen gradually. This is a generally beneficial feature of human decision making; negative emotional reactions to problems, such as fear and alarm, draw attention and motivate the individual to change their situation. However, the human brain seems to interpret the persistence of a problem as evidence that the problem is intractable, and hence not worthy of further attention, so it dials down the emotional response.

In some situations, and specifically when ongoing problems *do* warrant persistent efforts at mitigation, however, adaptation can lead to complacency and tolerance of what should be an intolerable situation. This is arguably the situation with privacy. In one of several experiments illustrating the importance of adaptation (Adjerid, Pe'er, & Acquisti, 2018), research participants were asked demographic questions, including their email address, were provided with specific privacy assurances, and then were asked various sensitive questions. Then, in a second round, they were provided with new privacy assurances and were asked a second set of sensitive questions. The experimental manipulations in the study involved whether the privacy assurances in the two rounds signaled increasing (low, then high), decreasing (high, then low), or stable (low, then low; or high, then high) privacy protections. The main result of the study was that the contrast between assurances mattered a lot. Information revelation was similar in the low–low and high–high conditions, suggesting that individuals did not have much of an independent idea of what levels of protection were appropriate. However, information revelation was much greater in the second round, in the low–high condition than in the high–low condition. Even in the short course of the experiment, subjects seemed to adapt to whatever level of privacy they were initially provided with, and became concerned when privacy assurances were reduced.

3.1.7 The Drive to Share

While much of the literature on privacy tends to focus on the risks of information leakage, those concerned about privacy need to contend with the fact that in most situations privacy-related motivations are counterpoised against potentially even stronger motives for socializing, connecting, and sharing information. Individuals have many reasons for sharing information, including economic benefits that result from strategic revelation or withholding, and an array of psychological motives (Carbone & Loewenstein, 2020). One study using subjective measures as well as fMRI (Tamir & Mitchell, 2012) found that information sharing is inherently pleasurable, particularly when the information relates to one's own thoughts and feelings.

Further research dating back decades documents health, psychological, and social benefits resulting from interpersonal disclosure (e.g., Pennebaker, 1997).

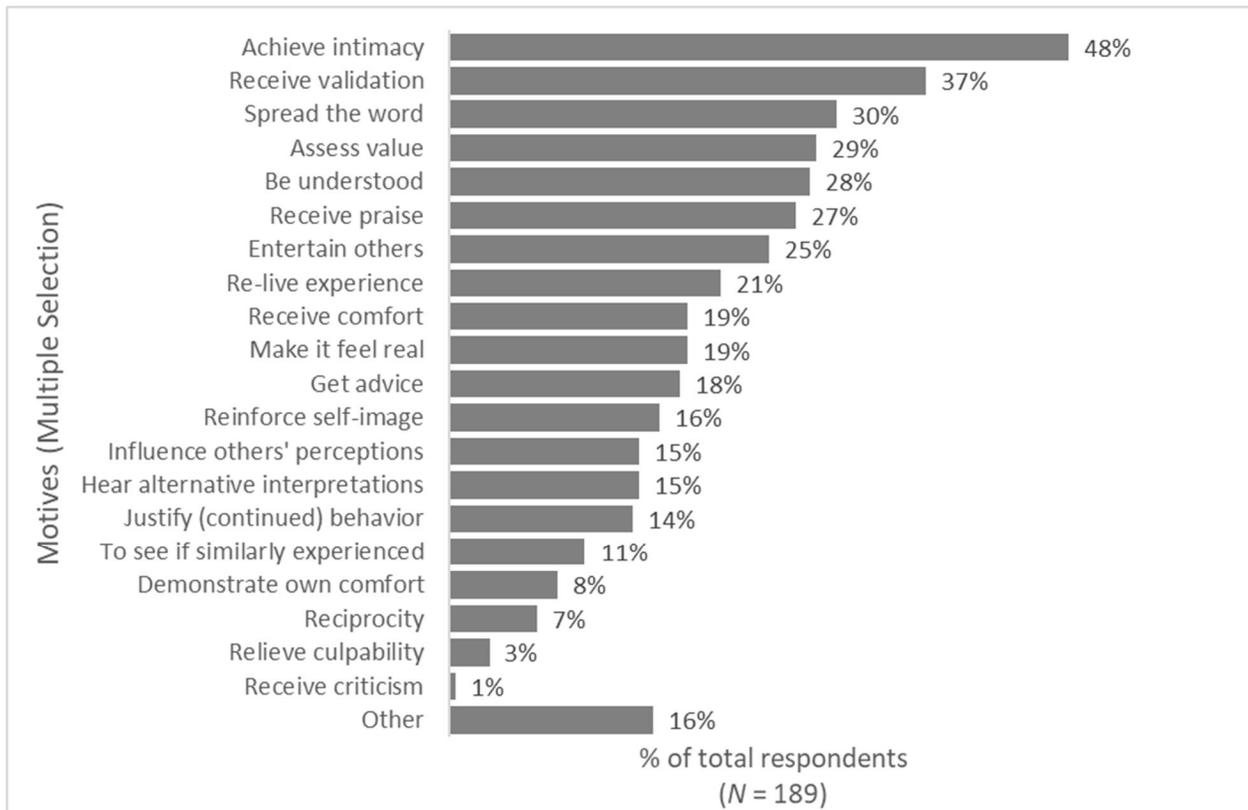


Figure 1 (Figure 6 from Carbone & Loewenstein, 2020)

Although the studies reported in Carbone and Loewenstein (2020) provide a range of insights about the desire to share information, perhaps the most relevant to the current paper is a finding regarding reasons for sharing. In that study, respondents to a survey (n=552) were asked an open-ended question about whether they could recall a situation in which they were “dying to share” information with another person; then, if they could recall such a situation, whether they had ended up sharing the information, as well as a series of follow-up questions about what information they were dying to share and who they were dying to share it with. Survey respondents were presented with a long list of possible motives and were asked to select all that they believe might have driven the intense desire to share their experiences.

Figure 1 presents the results from this analysis. The diversity of reasons that people cite to explain their desire to disclose underlines the power of motives counterpoised against those that drive privacy.

3.2 Implications

The research we have recounted implies that non-normative factors (factors that arguably do not affect true costs and benefits of sharing) can easily sway observed privacy choices. In theory, this may lead to outcomes that can either overshoot or underachieve desired levels of privacy. In reality, the ample survey evidence of widespread concerns about and desires for privacy (Section 2), as well as firms' deliberate use of those factors to nudge disclosure behavior (next, Section 4), indicate that—more often than not—online privacy behaviors may fall short of desires. In Altman's terms, achieved privacy does not meet desired privacy.

4. The Supply Side of Privacy

Psychological factors affect, and to some degree distort, observed market demand for privacy (Section 3). Economic factors affect its supply. In this section, we show that, due to the interaction between those two factors, even if consumers were infinitely savvy, they would still find desired (Section 4.1), as well as *desirable* (Section 4.2), levels of privacy nearly unattainable.

4.1 Privacy Under-supply

A supply of privacy does exist in the market. Techniques and protocols have been developed to protect data in nearly every imaginable online activity—from email to instant messaging, and from online advertising to online social media services. Some of those tools have been incorporated into products now available to consumers, such as VPN services, user-friendly encrypted emails, nontracking search engines and maps, anonymous browsers, and secure messaging platforms. Some major technology companies have started trying to leverage their private features as a source of competitive advantage (Panzarino, 2019). And most online

services offer to varying degrees privacy controls such as opt-outs, visibility settings, and so forth.

At the same time, living in the modern world means being subject to continuous and ubiquitous tracking. Whether we are aware of it or not, both our online and offline behaviors are constantly tracked, by surveillance cameras (Satariano, 2019), face-recognition technologies (Feng, 2019), rental cars activating GPS tracking by default (Mapon, 2017), multitudes of apps that share with an opaque ecosystem of intermediaries personal data from our phones (Almuhimedi et al., 2015), and trackers used by companies to learn and predict our browsing behavior (Federal Trade Commission, 2016). Even the boundaries between our online and offline existences are eroding: Statistical tools are used to match sets of data belonging to the same individual across domains and services, endangering the very notion of anonymity and permitting personal information to be inferred from nonsensitive public data (Narayanan & Shmatikov, 2008).

The forces that drive this relentless encroachment of surveillance into every facet of our lives are in part behavioral (our increasing adaptation and habituation to tracking) but in greater part economic: The reduction in the costs of both surveillance and data storage, as well as the increasing value of (and success in) monetizing personal data, increases firms' demand for tracking, thus driving down the supply of privacy. Such an outcome had been predicted by Hirshleifer (1978), who showed how private firms faced incentives to *overinvest* in gathering information: The resources used to acquire and disseminate that information would be wasteful from a societal point of view.

The rise of surveillance as an economic model (Zuboff 2015) has implications other than firms' overinvestments. First, markets with significant information asymmetries (such as, surely, the market for privacy) can lead to economic inefficiency and even market failures (Akerlof 1970). Furthermore, market forces lose, in part, their ability to restrain firms' data usage practices if network effects (especially powerful in two-sided market platforms such as search

engines, advertising networks, and social media) lead to quasi-monopoly incumbents. While the debate around competition, antitrust, and data regulation is nuanced and evolving, reduced competition and network effects reinforce each other, and enable incumbents to accumulate more data and user attention (say, from search behavior), then enter markets for other services (say, navigation maps), which in turn allows more data collection (and more user attention), and makes it possible for the incumbent to supply even more services than any entrant or competitor. Increased data collection—in terms of both increasing share of a consumer market and increasingly detailed inferences about each consumer—can thus lead to better, more valuable services, but also to fewer available outside options for privacy-seeking consumers.

Data-based network effects also tend to create powerful lock-ins into current products—such as online social networks, whose value resides precisely in continued engagement of a growing user base—strengthening incumbents. This creates what has been referred to as a privacy externality (Acquisti, Taylor, & Wagman, 2016): Other people's usage of privacy-intrusive services increases the cost for privacy-conscious consumers *not* to use them. At the extreme, consumers' costs to choose protective options become impractically large to accept. Once *no privacy* becomes the default, or the social norm, privacy options risk disappearing from the market altogether.

Making matters worse, consumers' *marginal* costs of protection increase rapidly with the degree of protection sought. Because so much of what we do is now tracked, there is simply too much to learn about and to protect. Everyone can, with little effort, refrain from publishing their social security number on their Facebook profile. Using a VPN is more costly—in cognitive, usability, and economic terms. Attempting to hide most of one's digital footprints from third-party monitoring is nearly incalculably demanding—and, in light of the continuously evolving technologies for surveillance, ultimately futile.

4.1.1 The Interaction of Economics and Psychology

While some firms may actively promote privacy, there is an almost surely greater fraction that respond to, and to a great extent exploit, consumers' psychological characteristics (Section 3) for their own ends. Before the term "dark patterns" started being popularized (Gray et al., 2018), behavioral research on privacy had highlighted how platform providers can leverage interface changes to influence privacy choice (Hartzog, 2010). Firms, for example, have a deep appreciation of the impact of defaults (Acquisti, John, & Loewenstein, 2013), and joining a long tradition (e.g., of enrolling consumers in magazine subscriptions that continue unless proactively terminated), they set privacy defaults in a liberal fashion, banking on consumers' laziness, and inattention (see Table 1 for a summary of different ways that firms can take advantage of consumer psychology with respect to privacy).

Even transparency and control can be used to nudge consumers toward higher disclosures (Acquisti, Adjerid, & Brandimarte, 2013). Transparency and control *are* important components of privacy management. For instance, they reduce customers' feeling of emotional violation and distrust in cases where personal data shared with firms are subject to vulnerabilities (Martin et al., 2017). But, for reasons expounded in Section 3, they are not *sufficient* conditions for privacy protection. In fact, economic factors (Section 4.1), such as network effects and lock-in, exacerbate behavioral hurdles (Section 3), giving some firms more data, more user attention, more control, and ultimately more ability to influence consumer behavior. Ultimately, consumer "responsibilization" (Giesler & Veresiu, 2014) approaches to privacy, predicated around so-called notice and consent regime—that is, reliance on consumer "choice" (Solove, 2012)—have not made it affordable or even feasible for consumers to achieve desired levels of privacy.

Consider the historical evolution of online tracking, which started with browser "cookies." As users started adopting cookie managers to prevent cross-site tracking, the online advertising industry developed so-called "flash" cookies to avoid consumers' deflecting strategies. And, as

users started acquiring tools to defend against this new form of tracking, the industry switched to even more intrusive—and harder to hide from—tracking strategies: device fingerprinting, deep packet inspection, and so on. The consumer-seeking privacy in the digital age cannot rely on Altman’s mutually shared social norms and intuitive behaviors, which worked in an offline world. She/he is a modern Sisyphus constantly forced to learn new strategies, to little avail.

4.2 Desired vs. *Desirable* Privacy

A valid counterpoint to the arguments in Section 4.1 is that desired levels of privacy, albeit unachievable, may in fact exceed what would actually be optimal for consumers. If, in the age of analytics, the collection and analysis of data can be source of great technological and economic advancement, then the loss of privacy, far from being a threat to societal well-being, may be a necessary price for increased consumer and societal welfare. Thus, *desirable* amounts of privacy may be less than what consumers claim to want. They may, in fact, be precisely the levels that markets already produce.

The fact that great utility can be extracted from data is undeniable and not questioned here. Rather, we scrutinize from a number of angles the premise that market outcomes satisfy optimal balances of data use and protection. What does privacy economics actually tell us about the trade-offs associated with data sharing (Section 4.2.1)? How are the benefits from data allocated (Section 4.2.2)? What are the societal costs of data protection, and can they be minimized? And finally, what are the costs of privacy invasions, and what dimensions of arguable concern for consumers are left out of economic analysis of data privacy (Section 4.2.3)?

4.2.1 Individual and Aggregate Consumer Trade-offs

A review of both theoretical research and empirical economic research on privacy (Acquisti, Taylor, & Wagman, 2016) belies the notion that more data collection is monotonically associated with positive changes in consumer welfare.

At the individual level, it is economically rational to want to decide for oneself what and how much to protect or reveal. It would be entirely logical (and welfare-increasing: Varian, 1996) for a consumer to share his/her product interests with marketers (to receive potentially beneficial targeted offers) and to hide the reservation price for those products (to avoid price discrimination). In a very Altmanian sense, the ability to regulate one's openness is consistent with economic arguments of utility maximization. The problem is that, as noted above, under current market conditions, there is no practical way for consumers to meaningfully regulate the use of their information. Since tracking is ubiquitous, once the ability to regulate data flows is taken away from consumers, secondary use of data is almost entirely beyond consumer control. And secondary use, in turn, can create significant negative externalities (Noam, 1997).

The same argument applies to the *aggregate* welfare effects of privacy protection (or lack thereof). First, the welfare implications of the collection and use of data are nuanced and context-dependent. A substantial body of modeling and empirical work, surveyed in Acquisti, Taylor, and Wagman (2016), finds that expansions in data collection and use are not necessarily welfare-increasing, and limitations not necessarily welfare-decreasing. Second, privacy trade-offs are inherently redistributive. Different decisions on data protection affect different stakeholders differently, and the interests of those stakeholders are rarely aligned (recall the consumer trying to hide his/her reservation price, and the seller trying to discern it). The theoretical goal of achieving desirable societal outcomes inevitably forces policymakers to tackle thorny questions regarding *whose* welfare they want to prioritize. Either by intervening with regulation, or by letting market outcomes determine levels of privacy across domains, they inescapably favor one or the other stakeholder.

4.2.2 Who Benefits from Consumer Data Collection?

Market equilibria may also not be advantageous for consumer welfare in relative terms, if most of the benefits accrued from the collection and usage of their data accrue to other

stakeholders, such as data intermediaries that, thanks to a combination of market power and supremacy in surveillance technology, have nearly unchallenged control over the data. Consider, for instance, online targeted advertising. According to industry insiders, collecting and using consumer data to target ads create economic win-wins for all parties involved—merchants, consumers, publishers, and intermediaries (AdExchanger, 2011). In reality, on theoretical grounds, targeting can either increase or decrease (Bergemann & Bonatti, 2011; De Corniere & De Nijs, 2016) the welfare of stakeholders other than the data intermediaries. And available *empirical* research tells us very little about how the benefits of targeted ads are allocated to stakeholders other than merchants and the intermediaries themselves. Data from ad networks suggest that opting out of behavioral targeting costs publishers and ad exchanges approximately \$8.5 per consumer (Johnson et al., 2020). Yet, a Digiday 2019 poll of publisher executives found that for 45% of respondents, behavioral ad targeting had “not produced any notable benefit,” and 23% claimed it had “actually caused their ad revenues to decline” (Weiss, 2019). And what about consumers themselves? Do consumer search costs go down because of targeted ads? Are the prices they pay for advertised products on average higher, or lower, than those they would have paid for products found via search? What about the quality of those products? Much more research is needed in this area.

4.2.3 The Costs of Privacy Protection

Two key approaches to privacy protection are regulation and privacy-enhancing technologies (PETs).

There is no shortage of studies showing privacy regulation to have, other than its intended impact of privacy protection itself, also negative effects on some economic metrics—for instance, loss of advertising effectiveness (Goldfarb & Tucker, 2011) or impaired technology adoption (Miller & Tucker, 2009). There is also evidence of privacy regulation having *positive* economic effects—for instance, reduction in identity theft (Romanosky et al., 2011) or increased technology adoption (Adjerid, Acquisti, et al., 2016). This seemingly contradictory evidence

should not be surprising: It is consistent with results from the broader economics literature, which show how the impact of regulation on innovation can be quite nuanced, and how the direction of the impact can vary depending on how interventions are designed, implemented, and enforced (BERR 2008). For instance, outcome- or performance-based interventions may score better than prescriptive regulations dictating the use of specific tools or technologies.

Similarly, in the privacy realm, we can reconcile contrasting findings by reiterating, first, the contextual nature of privacy trade-offs; second, by observing that binary metrics (such as absence vs. presence of regulation) are too coarse to capture the complex effects of privacy protection. The decisive factor in determining the effect of privacy regulation on innovation or welfare may not be its existence, but the *specifics* of the intervention (Goldfarb & Tucker, 2012). Well thought-out interventions can protect privacy and stimulate growth (Adjerid, Acquisti, et al., 2016); poorly thought-out ones may achieve neither goal. In addition, focusing on the short-term effects of regulatory intervention (as much empirical privacy work does for good reason—to identify robust causal links) risks missing not just the downstream beneficial effects of consumer protection, but also the long-term effects on competition and innovation arising from firms having to improve their services and systems to be more privacy-aware (for instance, by deploying privacy-enhancing technologies).

Several of such privacy-enhancing technologies (popularly referred to as PETs: Goldberg, 2002) use statistical and cryptographic techniques—such as differential privacy or homomorphic encryption—to allow extracting utility from data *while* protecting privacy. They demonstrate that intrusive data practices are not always essential to service effectiveness and that privacy protection is not inherently antithetical to data sharing and data analytics. As they can reduce data granularity, both regulation and PETs can certainly bear costs: Even mere privacy controls have been shown to reduce net contributions to crowdfunding campaigns (Burtch et al., 2015). But both computer science and economic research (Abowd & Schmutte,

2019) indicate that careful design can minimize those costs while ensuring some data protection.

4.2.4 Noneconomic Ramifications: Privacy “Dark Matter”

Our analysis has, so far, only focused on economically quantifiable implications of privacy choices in specific contexts—such as targeted advertising. Two critical considerations arise when we try to look at a broader picture.

First, costs of privacy across scenarios are arguably impossible to combine into a definitive, aggregate estimation of the “loss” caused by a generalized lack of privacy. And this is not because privacy costs are rare and few—but for the opposite reason: They are very common, but highly diverse in form, heterogeneous in likelihood, and varying in magnitude. They range from identity theft to price discrimination; from attention and time waste to psychological harm; from discrimination in targeted advertisement to filter bubbles; and from stigma to rare but catastrophic personal consequences (see, e.g., Acquisti, Taylor, & Wagman, 2016; Calo, 2011; Solove, 2005, 2007). Hence the aggregation and estimation problem. For instance, to the extent that data surreptitiously collected through privacy-intrusive apps can have an effect on a country’s election how can we quantify (or even demonstrate) that impact, and its plethora of downstream ramifications?

Second, and relatedly, we have not even considered many of the most consequential ramifications of the loss of privacy. We call this economic “dark matter”: We know it is there, but cannot quantify it. The ramifications include the collective value of privacy (Regan, 1995); its role in preserving room for individual subjectivity (Cohen, 2010); and its role in protecting freedom (Westin, 1967), dignity (Schoeman, 1984), and fairness (Jagadish, 2016), or the very integrity of social life (Nissenbaum, 2009). If market outcomes respond primarily to economic incentives, market equilibria may not account for these intricate, indirect, less tangible, and yet arguably even more critical implications of data protection.

4.2.5 Implications: Should Consumers Care?

To recap, market forces have driven an expanding encroachment of surveillance into consumers' lives over the past two decades (Section 4.1). This has been accompanied by undeniable economic benefits. And yet, economic analysis also raises valid doubts over the idea that increasing data collection necessarily enhances consumer or societal welfare (4.2.1), the notion of consumers accrue most of the value produced from their data (4.2.2), and the idea that data protection through regulation or technology will be inescapably welfare-decreasing (4.2.3). None of this discussion, moreover, recognizes the importance of noneconomic dimensions of privacy loss (4.2.4). We conclude that consumers have good reasons to aspire to higher degrees of privacy than market outcomes are likely to provide.

5. The Privacy “Paradox?”

The question of whether consumers *truly* care for privacy, the hurdles they face in achieving it, and the ramifications of their privacy choices in the market lie at the roots of the so-called “privacy paradox”—an apparent gap, or dichotomy, between people’s self-reported mental states (attitudes, concerns, desires, etc.) regarding privacy and their actual behaviors. A massive amount of scholarly effort over the past two decades has explored this gap, yet failed to affirmatively conclude whether a paradox of consumers claiming a desire for privacy, but not acting like they actually cared, does in fact exist (Norberg et al., 2007), or is a “myth” (Solove, 2021). Both the so-called paradox and the disagreements around it have profound implications for policy. Attempting to bring some clarity to this debate may help us better understand the policy ramifications of the research reviewed in the previous sections.

Evidence that some dichotomy (or, more simply, gaps of varying depth) between privacy mental states and actual behaviors *can* arise (but will not *always* arise) is actually common. Early observations focused on a broad version of the dichotomy: a gap between *generic* attitudes (typically expressing an overall desire for or concern over privacy) and actual behaviors (often interpreted as privacy-neglecting), such as social media sharing (Barnes, 2006)

or disclosure choices in experiments (Spiekermann et al., 2001). But inconsistencies between broad attitudes and specific behaviors are well-recognized in the literature predating privacy decision making (Ajzen & Fishbein, 1977); thus, considering the dynamic and dialectic nature of privacy boundary regulation, it would not be surprising to find them in the privacy realm as well (Dienlin & Trepte, 2015). And yet, evidence of specific, narrow dichotomies has also been often uncovered.

For instance, specific *attitudes* toward privacy in mobile phone apps were not closely linked to actual app download behavior in an incentivized experiment (Barth et al., 2019). Although participants were “concerned that mobile apps monitor[ed] their activities and that often too much personal information is collected” and “[t]he level of privacy intrusion by mobile apps in general was perceived to be relatively high,” nearly a third of participants downloaded the app that was ranked as the most intrusive within the set made available to the subjects, and 49% downloaded an app that was analyzed as intrusive—despite the fact that the participants had been provided extra funds to buy a nonintrusive app.

Specific *concerns* (captured via a questionnaire) did not match corresponding disclosure behaviors (captured from mined network data) for various members of the Carnegie Mellon Facebook network in 2006 (Acquisti & Gross, 2006). Nearly 16% of respondents who had expressed the highest degree of concern (on a 1- to 7-point Likert scale) for the scenario in which a stranger knew their schedule of classes and where they lived did in fact provide both pieces of information publicly, to everyone else on the network (around 7,000 users at the time). In fact, nearly 22% of those high-concern subjects provided their address, and nearly 40% provided their schedule of classes.

Mismatches between individuals’ specific *expectations* of privacy and actual sharing behaviors have also been documented. For instance, every single participant in a social media study exhibited some mismatch between what they believed their profiles’ privacy settings to be, and what they actually were (Madejski et al., 2012).

Also specific behavioral *intentions* have been shown to not necessarily match behaviors. Norberg et al. (2007) first asked individuals their intentions to disclose specific pieces of information to marketers, and then, several weeks later, asked the same subjects to actually provide “those exact pieces of information to a market researcher.” Self-reported intentions to disclose were significantly lower than actual disclosures. (Between-subject experiments - including some of those covered in the previous sections, in which non-normative factors are shown to significantly sway privacy choice - suggest similar mismatches: See, e.g., Adjerid, Pe’er, and Acquisti 2018.)

And yet, as noted, notwithstanding the evidence of gaps between privacy mental states and behaviors, the literature remains split between considering the paradox real and considering it a misnomer, or even a myth.

We believe these disagreements (and the state of confusion surrounding the so-called paradox) to be caused by two distinct factors.

The first factor is that different scholars have defined the dichotomy underlying the paradox in different ways (as already noted by Gerber et al., 2018 and Kokolakis, 2017): Different studies have focused on different privacy scenarios, as well as on different pairwise comparisons of mental states and behaviors (e.g., generic attitudes vs. specific intentions, or specific concerns vs. behaviors, or intentions vs. behaviors, and so on). Consequently, and considering the contextual nature of privacy choice, it should not be surprising that scholars found both evidence of correspondence between mental states and behaviors in some scenarios, and evidence of possible gaps (of varying depth) between mental states and behaviors in other scenarios. Sometimes evidence *for* and *against* a dichotomy can arise from within the same study. For instance, broad attitudes toward privacy were simultaneously found to be *uncorrelated* to disclosure patterns on Facebook, but correlated with the likelihood of joining the network (Acquisti & Gross, 2006).

The second (and perhaps more consequential) factor is that the very term “paradox” may have been interpreted differently by different scholars. “Veridical” paradoxes are those that include *seemingly* contradictory statements which, upon further observation, can actually be shown to be true (Quine 1976). We suspect that some privacy scholars, consistent with this definition, focused on the discovery of a seemingly contradictory gap between claimed privacy preferences and actual behaviors, and called that a “paradox” of privacy. Whereas other scholars focused on the *explanations* that resolved that apparent dichotomy, and since explanations did exist, concluded (contra Quine’s definition) that there was no paradox. In either case, the fact that misalignments between attitudes and behaviors can be explained does not imply that the misalignments themselves do not exist.

The debate over the existence of a paradox of privacy matters from a public policy perspective (Martin, 2020), as it reflects differing implications one could derive from the body of work on consumer privacy choice and market behaviors that we have reviewed in the previous sections. We highlight in the final part of this section the implications that seem to us more plausible.

First, even though there *is* evidence of situations in which behaviors do not match self-reported preferences, this does not imply that privacy behaviors never match preferences, or that self-reported preferences for privacy should never be trusted, or that consumers’ persistent demand for privacy in surveys should not be taken seriously. Those mismatches do not imply consumer carelessness or indifference for privacy; rather, they often have precise psychological and economic explanations, explored in previous sections. Second, by the same token, the fact that there is evidence of scenarios in which attitudes precisely predict behavior does not imply that consumers are *always* able to match their privacy desires with behaviors in the marketplace, and hence that no correcting public policy intervention is needed. Third, there is no *single* explanation that resolves the paradox. Rather, the explanations for the dichotomy – when a dichotomy does arise – are many, and not mutually exclusive, because many and diverse are

the factors affecting privacy choice, and the contexts in which the dichotomy can emerge (see discussions of multiple explanations in Acquisti, Brandimarte, & Loewenstein, 2015; Kokolakis, 2017; Solove, 2021). In fact, the explanations include nearly any of the psychological and economic factors we reviewed in Sections 3 and 4. This, in turn, implies that resolving the whole of the modern privacy problem amounts to much more than addressing just this or that specific concern.

Ultimately, the central lesson for policy of the paradox literature is that consumers may well care for privacy and try to regulate opening and closing of self to others in their everyday lives, but psychological and economic hurdles may make the privacy they desire unattainable in absence of a systemic, fundamental change in the way we approach the policy of privacy.

6. What Should Be Done?

If market outcomes are unlikely to produce not just the levels of privacy consumers desire, but also the levels that would be desirable for them, what—if anything—can be done to correct privacy imbalances?

6.1 Privacy Nudges

Some of the psychological hurdles we considered in Section 3 can be countered, or ameliorated, through behavioral interventions, to align privacy outcomes with ex ante preferences (Acquisti, 2009). Numerous privacy nudges have been explored in the literature, from changing social media default visibility settings to making the consequences of privacy choices more salient to users (Acquisti, Adjerid, Balebako, et al., 2017). Unfortunately, while nudges have been proven to be somewhat effective in experiments and field trials (for instance, Zhang & Xu, 2016), it is unclear that localized behavioral interventions alone can correct the enormous imbalance consumers encounter online between their ability to manage personal data and platform providers' ability to collect it. By controlling user interfaces, the providers remain in control of the architecture of choice.

6.2 Data as Property

Data propertization schemes have been proposed in the literature since the mid-1990s. Laudon (1996) proposed the creation of personal data markets, where consumers would trade rights with organizations over the collection and usage of their information, thereby “monetizing” their data. Over time, technological barriers to Laudon’s proposal have vanished. Data monetization startups have emerged and politicians have incorporated data propertization or “data dividends” in their platforms (Daniels, 2019). While appealing on some levels (Arrieta-Ibarra et al., 2018), data propertization schemes face hurdles in practice (Acquisti, Taylor, & Wagman, 2016). One issue is whether consumers, who under such a scheme will face decisions about who to sell their data to and for how much, are able to assign fair, reasonable valuations to their own data, considering the informational and behavioral hurdles we surveyed in Section 3. A second issue is that schemes that monetize privacy run the risk of exacerbating inequality, creating a world in which only the affluent can have privacy, or in which the already rich get more for their data than anyone else. Finally, considering the consequential noneconomic dimensions of privacy (Section 4.2.4), some might find abhorrent the notion of putting a price on it, or question the propriety of allowing people to sell it, much as many question whether people should be allowed to sell their own organs for transplant.

Furthermore, market-based data propertization schemes suffer from a nearly insurmountable economic challenge. The most valuable data are contextual and dynamic; it is *created* in the interaction between incumbent service providers and consumers; and—missing regulatory intervention establishing baseline protections—those providers are unlikely to relinquish its ownership to others. Hence, data propertization schemes may simply add themselves to an ecosystem of widespread surveillance, rather than replace it.

6.3 Privacy-Enhancing Technologies

Because they allow both data protection and data analytics, PETs offer significant potential individual and societal benefits. Yet, many consumers are unlikely to take advantage of them, due to unawareness of PETs’ existence, distrust, or perceived (and actual) costs. Thus,

barriers to the success of PETs are both psychological and economic in nature. Pushing the responsibility for their usage to individuals—that is, expecting them to navigate a universe of disparate, heterogeneous self-defense solutions across an ever-increasing range of scenarios in which data are tracked—would once again shift exorbitant costs onto consumers in usability, cognitive, and economic terms (such as the opportunity costs arising from loss of features in services when PETs are deployed). In any case, much as is the case for nudges and data propertization schemes, deployment of PETs is an inherently individualist solution: In the absence of a wide-ranging regulatory intervention supporting their deployment by making privacy the default (Cavoukian, 2009), it is hard to see how a patchwork approach of localized solutions—only working under specific circumstances, on specific apps or systems, in specific scenarios—could go far toward addressing what is inherently a *systemic* problem of privacy.

6.4 Privacy Regulation (and Its Economic and Behavioral Hurdles)

Psychologically informed interventions, data propertization schemes, and (especially) privacy-enhancing technologies may be useful tools for privacy management, but none is likely to work as intended in the absence of substantive, comprehensive policies that mandate a framework of baseline privacy protection, addressing both the consumer-side hurdles we considered in Section 3, and the supply-side factors we considered in Section 4. While we defer to the vast legal scholarship on privacy for a nuanced discussion on the effectiveness of different intervention models (such as legislation, litigation, and so forth), we consider here some psychological and economic factors affecting regulatory efforts.

Worldwide, there has been no shortage of privacy regulatory efforts. Recently, both in Europe (with the GDPR) and in the United States (with the California Consumer Privacy Act), major comprehensive initiatives aimed at addressing the novel digital privacy challenges have become law. Yet, despite the intentions of regulators to promote “privacy by design” (PBD) and “privacy by default” principles (which refer to organizations proactively considering privacy throughout the entire data lifecycle; see Cavoukian, 2009), a gulf still exists “between PBD in

principle and as implemented in practice” (Wong & Mulligan, 2019). The jury is still out on the effectiveness of recent regulatory interventions (Bamberger & Mulligan, 2019). For instance, a number of recent studies have suggested that, following the enactment of GDPR, consumer tracking remains ubiquitous (Sanchez-Rola et al., 2019).

Some hurdles impeding the enactment of comprehensive regulation are, again, psychological. A wide range of psychological factors can explain why there has not been an uprising of citizens to deal with the problem of privacy. Privacy has not yet become a “hot-button” issue (although some have argued that it could; see Taylor, 2001). Although in responses to closed-ended survey questions consumers report being concerned about privacy, this is rarely the case for open-ended questions that assess what problems are “top of mind” for respondents. For example, in the most recent Gallup poll in which voters were asked an open-ended question about the most important problem facing the nation, privacy (or any item that could be construed as related to privacy) did not make it to the list of the top 11 (the bottom of which garnered only 3% of mentions).

Why does the loss of privacy not generate, collectively, the kind of emotional response that some other societal problems do? The situation is reminiscent of the debate surrounding climate change, another great issue of our times. Certainly one reason is the lack of immediacy, and in some cases, of tangibility of the negative consequences of privacy invasions. Most people have, at best, a poorly formed notion of the negative consequences that can result from unregulated assaults on their data. In some cases, this is because they have not experienced them, and in other cases because they experienced them, but were unaware of it (e.g., when they pay higher prices for goods as a result of price discrimination enabled by sellers possessing their data).

Another, closely related reason is adaptation. People tend to adapt to—and pay little attention to—adverse situations that are either gradually changing or unchanging. This is especially true of situations that people feel powerless to change, so that the acceptance of

pronouncements such as “privacy is dead” is likely to make these self-fulfilling prophecies, introducing complacency about what has been lost and diminishing motivation for reform.

A second set of hurdles is economic in nature. “Regulatory capture” refers to the propensity for industries to play a central role in crafting the regulations that apply to them, resulting in a situation in which those regulations tend to favor the regulated industries—not the consumers that the regulations are ostensibly enacted to protect. Regulatory capture has been studied by economists (Downs, 1957), and more recently by privacy scholars (Hirsch, 2013; McGeeveran, 2016). According to this argument, since privacy regulation is in fact likely to come (and, in some cases, has come), companies will inevitably steer it in a direction that favors their interests over those of the consumers. By 2018, media were indeed reporting that major tech companies (including Apple, Google, and Facebook) *wanted* federal privacy regulation in the United States (Brandom, 2018).

A related concept from political economy has similar implications. It is the observation of political economist Mancur Olson (1965) that concentrated economic interests tend to trump diffuse, atomistic interests. Concentrated industries have incentives to lobby and influence policy in their favor. Individual citizens, in contrast, have a much more difficult time coordinating lobbying activity. Each individual, if only by dint of their limited resources, typically has only a limited interest in engaging in legislation-influencing tactics. There are certain exceptions to the pattern (e.g., hot-button issues such as guns and abortion, which can influence multitudes of people to join or form interest groups), but the advantage certainly resides with commercial, over the individual, interests.

The application of these ideas to privacy is straightforward. Several huge firms in the United States have both monumental resources (thanks to their ability to monetize consumer data) and incentives to lobby and influence public opinion in directions that propel privacy policies in their favor. In contrast, even the most privacy-conscious individual has limited ability, or likely motivation, to influence such policies. If privacy does not have the properties of hot-

button issues that lead people to rise up for reform, then privacy interest groups may not be able to contend with the organized lobbying of powerful firms.

7. Conclusion

The ultimate conclusion of this paper may appear pessimistic. We showed that people care and act to manage their privacy (Section 2), but face steep psychological (Section 3) and economic (Section 4) hurdles that make not just desired, but also desirable privacy nearly unattainable. We conclude that approaches to privacy management that rely purely on market forces and consumer responsabilization have failed. Comprehensive policy intervention is needed if a society's goal is to allow its citizens to be in the position to manage privacy effectively and to their best advantage. Yet, severe psychological and economic impediments to the enactment of regulation also exist.

More research—combining psychology, economics, computer science, and the law—is needed on those impediments, and how to overcome them. Research could take a nuanced perspective that examines both costs and benefits of different approaches to regulation, and that includes a consideration of trade-offs that are difficult to quantify, including the long-term ramifications of policy interventions. Research could also seek to assess how the benefits of consumer data collection are allocated to different stakeholders, and how different interventions can help to skew benefits toward the consumer side. Finally, research could address what types of privacy-protecting technologies would be most effective in enhancing consumer privacy, and most effective, in terms of producing benefits at minimum cost. Ultimately, targeted research could help scrutinize the premise that loss of privacy is unavoidable to enjoy the benefits of online services, and help uncover the best means for securing the best of both worlds.

Although our conclusions, as we noted, may appear pessimistic, some of the very evidence we discussed in this article provides a glimmer of hope. Turning back to where we started in Section 2, pronouncements that privacy is dead, we argued, confuse opportunities

with wants. People's opportunities for privacy are notably shrinking. And yet, across history, individuals—from Eastern Germans under Stasi (Betts, 2010; Sloan & Warner, 2016) to teenagers on social media (boyd & Marwick, 2011)—have revealed a remarkable tenacity in their attempts to carve out private spaces for themselves and their groups in face of all odds—even while in public, and even under surveillance. Technologies, interfaces, and market forces can all influence human behavior. But probably, and hopefully, they cannot alter human nature. If privacy, as Altman proposed, is both culturally specific and culturally universal, chances are that people's quest for privacy will not dissipate.

References

- Abowd, J. M., & Schmutte, I. M. (2019). An economic analysis of privacy protection and statistical accuracy as social choices. *American Economic Review*, 109(1), 171–202.
- Acquisti, A. (2004). Privacy in electronic commerce and the economics of immediate gratification. In *Proceedings of the 5th ACM conference on electronic commerce*.
- Acquisti, A. (2009). Nudging privacy: The behavioral economics of personal information. *IEEE Security & Privacy*, 7(6), 82–85.
- Acquisti, A., Adjerid, I., Balebako, R., Brandimarte, L., Cranor, L.F., Komanduri, S., Leon, P.G., Sadeh, N., Schaub, F., Sleeper, M., & Wang, Y. (2017). Nudges for privacy and security: Understanding and assisting users' choices online. *ACM Computing Surveys (CSUR)*, 50(3), 1–41.
- Acquisti, A., Adjerid, I., & Brandimarte, L. (2013). Gone in 15 seconds: The limits of privacy transparency and control. *IEEE Security & Privacy*, 11(4), 72–74.
- Acquisti, A., Brandimarte, L., & Loewenstein, G. (2015). Privacy and human behavior in the age of information. *Science*, 347(6221), 509–514.
- Acquisti, A., & Gross, R. (2006). Imagined communities: Awareness, information sharing, and privacy on the Facebook. In *International workshop on privacy enhancing technologies*. Springer.
- Acquisti, A., John, L. K., & Loewenstein, G. (2012). The impact of relative standards on the propensity to disclose. *Journal of Marketing Research*, 49(2), 160–174.
- Acquisti, A., John, L. K., & Loewenstein, G. (2013). What is privacy worth? *The Journal of Legal Studies*, 42(2), 249–274.

- Acquisti, A., Taylor, C., & Wagman, L. (2016). The economics of privacy. *Journal of Economic Literature*, 54(2), 442–492.
- Adjerid, I., Acquisti, A., Telang, R., Padman, R., & Adler-Milstein, J. (2016). The impact of privacy regulation and technology incentives: The case of health information exchanges. *Management Science*, 62(4), 10421063.
- AdExchanger. (2011, October 28). If a consumer asked you, “Why is tracking good?”, what would you say? <https://adexchanger.com/online-advertising/why-is-tracking-good/>
- Adjerid, I., Pe’er E., & Acquisti, A. (2018). Beyond the privacy paradox: Objective versus relative risk in privacy decision making. *MIS Quarterly*, 42(2), 465–488.
- Ajzen, I., & Fishbein, M. (1977). Attitude-behavior relations: A theoretical analysis and review of empirical research. *Psychological Bulletin*, 84(5), 888–918.
- Akerlof, G. A. (1970). The Market for "Lemons": Quality Uncertainty and the Market Mechanism. *The Quarterly Journal of Economics*, 84(3), 488-500.
- Almuhimedi, H., Schaub, F., Sadeh, N., Adjerid, I., Acquisti, A., Gluck, J., Cranor, L., & Agarwal, Y. (2015, April). Your location has been shared 5,398 times! A field study on mobile app privacy nudging. In *Proceedings of the 33rd annual ACM Conference on Human Factors in Computing Systems* (pp. 787–796).
- Altman, I. (1975). *The environment and social behavior*. Brooks/Cole Pub. Co.
- Altman, I. (1977). Privacy regulation: Culturally universal or culturally specific? *Journal of Social Issues*, 33(3), 66–84.
- Arrieta-Ibarra, I., Goff, L., Jiménez-Hernández, D., Lanier, J., & Weyl, E. G. (2018). Should we treat data as labor? Moving beyond "free". *AEA Papers and Proceedings*, 108, 38–42.
- Athey, S., Catalini, C., & Tucker, C. (2017). *The digital privacy paradox: Small money, small costs, small talk* (No. w23488). National Bureau of Economic Research.
- Bamberger, K. A., & Mulligan, D. K. (2019). Privacy law: On the books and on the ground. In *The handbook of privacy studies: An interdisciplinary introduction* (p. 349).
- Barassi, V. (2019). Datafied citizens in the age of coerced digital participation. *Sociological Research Online*, 24(3), 414–429.
- Barnes, S. B. (2006). A privacy paradox: Social networking in the United States. *First Monday*, 11(9).
- Barth, S., de Jong, M. D., Junger, M., Hartel, P. H., & Roppelt, J. C. (2019). Putting the privacy paradox to the test: Online privacy and security behaviors among users with technical knowledge, privacy awareness, and financial resources. *Telematics and informatics*, 41, 55–69.
- Bergemann, D., & Bonatti, A. (2011). Targeting in advertising markets: Implications for offline versus online media. *The RAND Journal of Economics*, 42(3), 417–443.

Beresford, A. R., Kübler, D., & Preibusch, S. (2012). Unwillingness to pay for privacy: A field experiment. *Economics Letters*, 117(1), 25–27.

BERR (Department for Business, Enterprise, and Regulatory Reform) (2008). Regulation and innovation: evidence and policy implications. *BERR Economics Paper n. 4*, United Kingdom.

Bettinger, E. P., Long, B. T., Oreopoulos, P., & Sanbonmatsu, L. (2012). The role of application assistance and information in college decisions: Results from the H&R Block FAFSA experiment. *The Quarterly Journal of Economics*, 127(3), 1205–1242.

Betts, P. (2010). *Within walls: Private life in the German Democratic Republic*. Oxford University Press.

Boyd, D., & Marwick, A. E. (2011). Social privacy in networked publics: Teens' attitudes, practices, and strategies. In *A decade in Internet time: Symposium on the dynamics of the Internet and society*.

Brandimarte, L., Acquisti, A., & Loewenstein, G. (2013). Misplaced confidences: Privacy and the control paradox. *Social Psychological and Personality Science*, 4(3), 340–347.

Brandom, R. (2018, October 24). Tim Cook wants a federal privacy law—But so do Facebook and Google. *The Verge*. <https://www.theverge.com/2018/10/24/18018686/tim-cook-apple-privacy-law-facebook-google-gdpr>

Burtch, G., Ghose, A., & Wattal, S. (2015). The hidden cost of accommodating crowdfunder privacy preferences: A randomized field experiment. *Management Science*, 61(5), 949–962.

Bush, G. W. (2009). Economic regulation. Chapter 9. White House Archives. https://georgewbush-whitehouse.archives.gov/cea/ERP_2009_Ch9.pdf

Calo, R. (2011). The boundaries of privacy harm. *Indiana Law Journal*, 86, 1131.

Carbone, E., & Loewenstein, G. (2020.) Dying to divulge: The determinants of, and relationship between, desired and actual disclosure. <https://ssrn.com/abstract=3613232>.

Cavoukian, A. (2009). *Privacy by design: The 7 foundational principles*. Information and Privacy Commissioner of Ontario, Canada.

Chen, B. X. (2018, March 21). Want to #DeleteFacebook? You can try. *New York Times*. <https://www.nytimes.com/2018/03/21/technology/personaltech/delete-facebook.html>

CISCO (2019). Consumer Privacy Survey: The growing imperative of getting data privacy right. CISCO Cybersecurity Series 2019. <https://www.cisco.com/c/dam/en/us/products/collateral/security/cybersecurity-series-2019-cps.pdf>

Cohen, J. E. (2010). What privacy is for. *Harvard Law Review*, 126, 1904.

Daniels, J. (2019, February 12). California governor proposes 'new data dividend' that could call on Facebook and Google to pay users. *CNBC*. <https://www.cnbc.com/2019/02/12/california-gov-newsom-calls-for-new-data-dividend-for-consumers.html>

De Corniere, A., & De Nijs, R. (2016). Online advertising and privacy. *The RAND Journal of Economics*, 47(1), 48–72.

Dienlin, T., & Trepte, S. (2015). Is the privacy paradox a relic of the past? An in-depth analysis of privacy attitudes and privacy behaviors. *European Journal of Social Psychology*, 45(3), 285–297.

Downs, A. (1957). *An economic theory of democracy*. Harper & Row.

Draper, N. A., & Turow, J. (2019). The corporate cultivation of digital resignation. *New Media & Society*, 21(8), 1824–1839.

DuckDuckGo. (2017, January). *A study on private browsing: Consumer usage, knowledge, and thoughts*. Technical report. https://duckduckgo.com/download/Private_Browsing.pdf

Federal Trade Commission. (2016, June). Online tracking. <https://www.consumer.ftc.gov/articles/0042-online-tracking>

Feng, E. (2019, December 16). How China is using facial recognition technology. *NPR*. <https://www.npr.org/2019/12/16/788597818/how-china-is-using-facial-recognition-technology>

Fiegerman, S. (2017, September 7). The biggest data breaches ever. *CNN Business*. <http://money.cnn.com/2017/09/07/technology/business/biggest-breaches-ever/index.html>

Fiesler, C., Dye, M., Feuston, J. L., Hiruncharoenvate, C., Hutto, C. J., Morrison, S., Roshan, P. K., Pavalanathan, U., Bruckman, A. S., De Choudhury, M., & Gilbert, E. (2017). What (or who) is public? Privacy settings and social media content sharing. In *Proceedings of the 2017 ACM conference on computer supported cooperative work and social computing* (pp. 567–580).

Gerber, N., Gerber, P., & Volkamer, M. (2018). Explaining the privacy paradox: A systematic review of literature investigating privacy attitude and behavior. *Computers & Security*, 77, 226–261.

Ghose, A. (2017). *Tap: Unlocking the mobile economy*. MIT Press.

Giesler, M., & Veresiu, E. (2014). Creating the responsible consumer: Moralistic governance regimes and consumer subjectivity. *Journal of Consumer Research*, 41(3), 840–857.

Godinho de Matos, M., & Adjerid, I. (2019). *Consumer behavior and firm targeting after GDPR: The case of a telecom provider in Europe*. NBER Summer Institute on IT and Digitization.

Goldberg, I. (2002). Privacy-enhancing technologies for the Internet, II: Five years later. In *International workshop on privacy enhancing technologies*. Springer.

Goldfarb, A., & Tucker, C. E. (2011). Privacy regulation and online advertising. *Management Science*, 57(1), 57–71.

- Goldfarb, A., & Tucker, C. (2012). Privacy and innovation. *Innovation Policy and the Economy*, 12(1), 65–90.
- Gray, C. M., Kou, Y., Battles, B., Hoggatt, J., & Toombs, A. L. (2018). The dark (patterns) side of UX design. In *Proceedings of the 2018 CHI conference on human factors in computing systems* (pp. 1–14).
- Gross, R., & Acquisti, A. (2005). Information revelation and privacy in online social networks (The Facebook case). In *Proceedings of the 2005 ACM workshop on privacy in the electronic society* (pp. 71–80).
- Habib, H., Colnago, J., Gopalakrishnan, V., Pearman, S., Thomas, J., Acquisti, A., Christin, N., & Cranor, L. F. (2018). Away from prying eyes: Analyzing usage and understanding of private browsing. In *Fourteenth symposium on usable privacy and security (SOUPS 2018)* (pp. 159–175).
- Hartzog, W. (2010). Website design as contract. *American University Law Review*, 60, 1635.
- Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The weirdest people in the world?. *Behavioral and Brain Sciences*, 33(2–3), 61–83.
- Hirsch, D. D. (2013). Going Dutch: Collaborative Dutch privacy regulation and the lessons it holds for US privacy law. *Michigan State Law Review*, 83.
- Hirshleifer, J. (1978). The private and social value of information and the reward to inventive activity. In *Uncertainty in economics* (pp. 541–556). Academic Press.
- Hoofnagle, C. J., & Urban, J. M. (2014). Alan Westin's privacy homo economicus. *Wake Forest Law Review*, 49, 261.
- Jagadish, H. V. (2016). The values challenge for Big Data. In *Bulletin of the IEEE computer society technical committee on data engineering* (pp. 77–84).
- Jentzsch, N., Preibusch, S., & Harasser, A. (2012). *Study on monetising privacy. an economic model for pricing personal information*. European Network and information Security Agency (ENISA).
- John, L., Acquisti, A., & Loewenstein, G. (2011). Strangers on a plane: Context-dependent willingness to divulge sensitive information. *Journal of Consumer Research*, 37(5), 858–873.
- Johnson, B. (2020, January 11). Privacy no longer a social norm, says Facebook founder. *he Guardian*. <https://www.theguardian.com/technology/2010/jan/11/facebook-privacy>
- Johnson, G. A., Shriver, S. K., & Du, S. (2020). Consumer privacy choice in online advertising: Who opts out and at what cost to industry?. *Marketing Science*, 39(1), 33–51.
- Kang, R., Dabbish, L., Fruchter, N., & Kiesler, S. (2015). “My data just goes everywhere”: User mental models of the Internet and implications for privacy and security. In *Eleventh symposium on usable privacy and security (SOUPS 2015)* (pp. 39–52).

- KFF. (2020, April). *Coronavirus, social distancing, and contact tracing*. Health tracking poll. <https://www.kff.org/global-health-policy/issue-brief/kff-health-tracking-poll-late-april-2020/>
- Kokolakis, S. (2017). Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon. *Computers & Security*, 64, 122–134.
- Kunreuther, H., Ginsberg, R., Miller, L., Sagi, P., Slovic, P., Borkan, B., & Katz, N. (1978). *Disaster insurance protection: Public policy lessons*. Wiley.
- Laudon, K. C. (1996). Markets and privacy. *Communications of the ACM*, 39(9), 92–104.
- Lewis, B. (2017, November 7). Americans Say Data Privacy is Important, but Few Take Steps to Protect Themselves. Instamotor. <https://instamotor.com/blog/online-data-privacy-survey>
- Madejski, M., Johnson, M., & Bellovin, S. M. (2012, March). A study of privacy settings errors in an online social network. In *2012 IEEE international conference on pervasive computing and communications workshops* (pp. 340–345). IEEE.
- Mapon. (2017, June 9). GPS tracking for rental cars: How to break from the mold. <https://www.mapon.com/us-en/blog/2017/06/gps-tracking-for-rental-cars-how-to-break-from-the-mold>
- Margulis, S. T. (2003). Privacy as a social issue and behavioral concept. *Journal of social issues*, 59(2), 243-261.
- Marreiros, H., Tonin, M., Vlassopoulos, M., & Schraefel, M. C. (2017). “Now that you mention it”: A survey experiment on information, inattention and online privacy. *Journal of Economic Behavior & Organization*, 140, 1–17.
- Martin, K. (2020). Breaking the privacy paradox: the value of privacy and associated duty of firms. *Business Ethics Quarterly*, 30(1), 65–96.
- Martin, K. D., Borah, A., & Palmatier, R. W. (2017). Data privacy: Effects on customer and firm performance. *Journal of Marketing*, 81(1), 36–58.
- Martin, K., & Nissenbaum, H. (2016). Measuring privacy: An empirical test using context to expose confounding variables. *Columbia Science and Technology Law Review*, 18(1), 176–218.
- McDonald, A. M., & Cranor, L. F. (2008). The cost of reading privacy policies. *Information System: A Journal of Law and Policy for the Information Society*, 4, 543.
- McGeveran, W. (2016). Friending the privacy regulators. *Arizona Law Review*, 58, 959.
- Melumad, S., & Meyer, R. (2020). Full disclosure: How smartphones enhance consumer self-disclosure. *Journal of Marketing*, 84(3), 28–45.

- Miller, C. C. (2014, November 12). Americans say they want privacy, but act as if they don't. *New York Times*. <https://www.nytimes.com/2014/11/13/upshot/americans-say-they-want-privacy-but-act-as-if-they-dont.html>
- Miller, A. R., & Tucker, C. (2009). Privacy protection and technology diffusion: The case of electronic medical records. *Management Science*, 55(7), 1077–1093.
- Moore Jr, B. (1984). *Privacy: Studies in social and cultural history*. Routledge.
- Murphy, R. F. (1964). Social distance and the veil. *American Anthropologist*, 66(6), 1257–1274.
- Narayanan, A., & Shmatikov, V. (2008). Robust de-anonymization of large sparse datasets. In *2008 IEEE symposium on security and privacy*. IEEE.
- Nissenbaum, H. (2009). *Privacy in context: Technology, policy, and the integrity of social life*. Stanford University Press.
- Noam, E. M. (1997). Privacy and self-regulation: Markets for electronic privacy. *Privacy and Self-Regulation in the Information Age*, 21–33.
- Norberg, P. A., Horne, D. R., & Horne, D. A. (2007). The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of Consumer Affairs*, 41(1), 100–126.
- Olson, M. (1965). *The Logic of Collective Action*. Cambridge University Press.
- Palen, L., & Dourish, P. (2003). Unpacking" privacy" for a networked world. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 129-136).
- Panzarino, M. (2019, March 19). Apple ad focuses on iPhone's most marketable feature: Privacy. *Techcrunch*. <https://techcrunch.com/2019/03/14/apple-ad-focuses-on-iphones-most-marketable-feature-privacy/>
- Pennebaker, J. W. (1997). *Opening up: The healing power of emotional expression*. Guilford.
- Penney, J. W. (2016). Chilling effects: Online surveillance and Wikipedia use. *Berkeley Technology Law Journal*, 31, 117.
- Petronio, S. (2002). *Boundaries of privacy: Dialectics of disclosure*. Suny Press.
- Pew Research Center. (2012). *Privacy management on social media sites*. <https://www.pewresearch.org/internet/2012/02/24/privacy-management-on-social-media-sites/>.
- Pew Research Center. (2013). *Anonymity, privacy and security online*. <https://www.pewresearch.org/internet/2013/09/05/anonymity-privacy-and-security-online/>.
- Pew Research Center. (2015). *Americans' privacy strategies post-Snowden*. https://www.pewresearch.org/wp-content/uploads/sites/9/2015/03/PI_AmericansPrivacyStrategies_0316151.pdf.

Pew Research Center. (2019). *Americans and privacy: Concerned, confused and feeling lack of control over their personal information*. https://www.pewresearch.org/internet/wp-content/uploads/sites/9/2019/11/Pew-Research-Center_PI_2019.11.15_Privacy_FINAL.pdf.

Posner, R. A. (1978). Economic theory of privacy. *Regulation*, 2, 19–26.

Preibusch, S., Kübler, D., & Beresford, A. R. (2013). Price versus privacy: An experiment into the competitive advantage of collecting less personal information. *Electronic Commerce Research*, 13(4), 423–455.

Quine, W. V. (1976). *The ways of paradox*. Harvard University Press.

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

Romanosky, S., Telang, R., & Acquisti, A. (2011). Do data breach disclosure laws reduce identity theft? *Journal of Policy Analysis and Management* 30(2), 256–286.

Rosenberg, M., Confessore, N., & Cadwalladr, C. (2018, March 17). How Trump consultants exploited the Facebook data of millions. *New York Times*.
<https://www.nytimes.com/2018/03/17/us/politics/cambridge-analytica-trump-campaign.html>

Sanchez-Rola, I., Dell'Amico, M., Kotzias, P., Balzarotti, D., Bilge, L., Vervier, P. A., & Santos, I. (2019, July). Can I opt out yet? GDPR and the global illusion of cookie control. In *Proceedings of the 2019 ACM Asia conference on computer and communications security* (pp. 340–351).

Satariano, A. (2019, September 15). Real-time surveillance will test the British tolerance for cameras. *New York Times*. <https://www.nytimes.com/2019/09/15/technology/britain-surveillance-privacy.html>

Savage, S. J. & Waldman, D. M. (2015). Privacy tradeoffs in smartphone applications. *Economics Letters*, 137, 171–175.

Schoeman, F. (1984). Privacy: Philosophical Dimensions. *American Philosophical Quarterly*, 21(3), 199-213.

Sloan, R. H., & Warner, R. (2016). The self, the Stasi, and NSA: Privacy, knowledge, and complicity in the surveillance state. *Minnesota Journal of Law, Science and Technology*, 17, 347.

Slovic, P. (1995). The construction of preference. *American Psychologist*, 50(5), 364.

Solove, D. J. (2005). A taxonomy of privacy. *University of Pennsylvania Law Review*, 154, 477.

Solove, D. J. (2007). I've got nothing to hide and other misunderstandings of privacy. *San Diego Law Review*, 44, 745.

Solove, D. J. (2012). Introduction: Privacy self-management and the consent dilemma. *Harvard Law Review*, 126, 1880.

- Solove, D. (2021, forthcoming). The myth of the privacy paradox. *George Washington Law Review*, 89.
- Spiekermann, S., Grossklags, J., & Berendt, B. (2001, October). E-privacy in 2nd generation e-commerce: Privacy preferences versus actual behavior. In *Proceedings of the 3rd ACM conference on electronic commerce* (pp. 38–47).
- Sprenger, P. (1999, January 26). Sun on privacy: “Get over it.” *Wired*. <https://www.wired.com/1999/01/sun-on-privacy-get-over-it/>
- Statista. (2016, November 23). Number of mobile phone users worldwide from 2015 to 2020. <https://www.statista.com/statistics/274774/forecast-of-mobile-phone-users-worldwide/>
- Stutzman, F. D., Gross, R., & Acquisti, A. (2013). Silent listeners: The evolution of privacy and disclosure on Facebook. *Journal of Privacy and Confidentiality*, 4(2), 7–41.
- Svirsky, D. (2019). *Three experiments about human behavior and legal regulation* [Doctoral dissertation, Harvard University, Graduate School of Arts & Sciences].
- Taylor, H. (2001). Testimony on “Opinion surveys: What consumers have to say about information privacy.” Hearing before the Subcommittee on Commerce, Trade and Consumer Protection. Serial No. 107-35. <https://www.govinfo.gov/content/pkg/CHRG-107hrg72825/html/CHRG-107hrg72825.htm>
- Tamir, D. I., & Mitchell, J. P. (2012). Disclosing information about the self is intrinsically rewarding. *Proceedings of the National Academy of Sciences*, 109(21), 8038–8043.
- Tsai, J. Y., Egelman, S., Cranor, L., & Acquisti, A. (2011). The effect of online privacy information on purchasing behavior: An experimental study. *Information Systems Research*, 22(2), 254–268.
- Varian, H. R. (1996). Economic aspects of personal privacy. In *Privacy and self-regulation in the information age*. National Telecommunications and Information Administration, US Department of Commerce.
- Vitak, J., & Ellison, N. B. (2013). “There’s a network out there you might as well tap”: Exploring the benefits of and barriers to exchanging informational and support-based resources on Facebook. *New Media & Society*, 15(2), 243–259.
- Vitak, J., & Kim, J. (2014). “You can’t block people offline”: Examining how Facebook’s affordances shape the disclosure process. In *Proceedings of the 17th ACM conference on computer supported cooperative work & social computing* (pp. 461–474).
- Weiss, M. (2019, June 5). Digiday research: Most publishers don’t benefit from behavioral ad targeting. *Digiday*. <https://digiday.com/media/digiday-research-most-publishers-dont-benefit-from-behavioral-ad-targeting/>
- Westin, A. (1967). *Privacy and freedom*. Atheneum.

Westin, A. (2001). Testimony on “Opinion surveys: What consumers have to say about information privacy.” Hearing before the Subcommittee on Commerce, Trade and Consumer Protection. Serial No. 107-35. <https://www.govinfo.gov/content/pkg/CHRG-107hrg72825/html/CHRG-107hrg72825.htm>

White, T. B., Novak, T. P., & Hoffman, D. L. (2014). No strings attached: When giving it away versus making them pay reduces consumer information disclosure. *Journal of Interactive Marketing, 28*(3), 184–195.

Wong, J. C. (2019, March 18). The Cambridge Analytica scandal changed the world: But it didn't change Facebook. *The Guardian*.
<https://www.theguardian.com/technology/2019/mar/17/the-cambridge-analytica-scandal-changed-the-world-but-it-didnt-change-facebook>

Wong, R. Y., & Mulligan, D. K. (2019). Bringing design to the privacy table: Broadening “design” in “privacy by design” through the lens of HCI. In *Proceedings of the 2019 CHI conference on human factors in computing systems*.

Zhang, B. & Xu, H. (2016). Privacy nudges for mobile applications: Effects on the creepiness emotion and privacy attitudes. In *Proceedings of the 19th ACM conference on computer-supported cooperative work & social computing*.

Zuboff, S. (2015). Big other: surveillance capitalism and the prospects of an information civilization. *Journal of Information Technology, 30*(1), 75-89.