What Is Privacy Worth?

Citation

Acquisti, Alessandro, Leslie K. John, and George Loewenstein. "What Is Privacy Worth?" *Journal of Legal Studies* 42, no. 2 (June 2013): 249–274. doi:10.1086/671754.

Published Version

doi:10.1086/671754

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:37101490>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#OAP>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

What Is Privacy Worth?

Alessandro Acquisti, Leslie K. John, and George Loewenstein

ABSTRACT

Understanding the value that individuals assign to the protection of their personal data is of great importance for business, law, and public policy. We use a field experiment informed by behavioral economics and decision research to investigate individual privacy valuations and find evidence of endowment and order effects. Individuals assigned markedly different values to the privacy of their data depending on (1) whether they were asked to consider how much money they would accept to disclose otherwise private information or how much they would pay to protect otherwise public information and (2) the order in which they considered different offers for their data. The gap between such values is large compared with that observed in comparable studies of consumer goods. The results highlight the sensitivity of privacy valuations to contextual, nonnormative factors.

1. INTRODUCTION

Understanding the value that individuals assign to the protection of their personal data is of great importance to businesses, the legal community, and policy makers. It is important to businesses because by estimating how much customers value the protection of their personal data, managers can seek to predict which privacy-enhancing initiatives may become sources of competitive advantage and which intrusive initiatives may trigger adverse reactions.

It is important to legal scholars and practitioners because privacy is an issue that has become increasingly prominent in the law in recent years, in part because of the emergence of new technologies, such as tracking by global positioning system (GPS) and social networking over

ALESSANDRO ACQUISTI is Associate Professor of Information Technology and Public Policy, Carnegie Mellon University. LESLIE K. JOHN is Assistant Professor of Business Administration, Harvard Business School. GEORGE LOEWENSTEIN is the Herbert A. Simon Professor of Economics and Psychology, Carnegie Mellon University.

[*Journal of Legal Studies*, vol. 42 (June 2013)]

© 2013 by The University of Chicago. All rights reserved 0047-2530/2013/4202-0009\$10.00

the Internet. In a recent case described in the *Washington Post* (Barnes 2012, p. A1) as “a first test of how privacy rights will be protected in the digital age,” the Supreme Court unanimously overturned the conviction and lifetime sentence of a Washington, D.C., drug dealer on the basis of the argument that monitoring the movements of his Jeep by affixing a GPS device to it for 28 days violated his Fourth Amendment rights (see *United States v. Antoine Jones*, 132 S. Ct. 945 [2012]). As has often been pointed out, the U.S. Constitution does not contain any explicit protection of privacy, so the judiciary has been searching for ways of connecting existing constitutional protections, such as the Fourth Amendment’s protection against unreasonable search and seizure, with the privacy issues of the day. In navigating the complex issues of privacy and attempting to reach a desirable balance between the goals of information sharing and commerce, on the one hand, and protection of personal information, on the other, the judiciary has sometimes sought guidance from estimates of the valuations that people assign to their privacy (Romanosky and Acquisti 2009).

Finally, individual valuations of privacy are important to policy makers, who are often required to choose between policies that trade off privacy for other desirable goals. For example, the Health Insurance Portability and Accountability Act of 1996 (HIPAA) gave patients greater privacy protections than they previously had but at the price of increased administrative costs and bureaucracy. Whether the changes wrought by HIPAA are worth their cost depends, at least in part, on the value that people place on privacy.

In recent years, there has been no shortage of empirical studies attempting to quantify individual privacy valuations in diverse contexts—such as personal information revealed online (Hann et al. 2007), location data (Cvrcek et al. 2006), or removal from marketers’ call lists (Varian, Wallenberg, and Woroch 2005). Some of these studies—as well as anecdotal evidence about the popularity of blogs, online social networks, and other information-sharing social media—suggest that even ostensibly privacy-conscious individuals are likely to share sensitive information with strangers (Spiekermann, Grossklags, and Berendt 2001). Applying the economic principle of revealed preferences, some have concluded that our society, quite simply, does not place much value on privacy (Rubin and Lenard 2002). Is it really possible, however, to measure the value that people place on privacy? And has less privacy truly become the new social norm, as a prominent Web 2.0 chief executive officer has argued (Gonsalves 2010)?

Another key aspect of policy concerns the degree to which issues of privacy warrant regulation. In the aftermath of a spate of well-publicized data breaches and identity thefts, U.S. legislators have introduced bills to regulate how businesses collect and protect consumer information,¹ and regulators have published guidelines and best practices for consumer data protection (U.S. Department of Commerce 2010; U.S. Federal Trade Commission 2010). However, whether regulators and legislators should intervene in the market for privacy is heavily debated in legal (Solove 2004) and economic (Lenard and Rubin 2010) circles. Some writers have proposed that U.S. policy makers should rely on self-regulatory frameworks (Lenard and Rubin 2010), which are predicated on the assumption that consumers can form sensible valuations of their personal information (and of the costs that arise when that information is compromised) and respond in a calculated, rational, fashion—an assumption that has not, so far, received empirical support.

The roots of economic research on privacy (which can be found in seminal writings of scholars such as Richard Posner and George Stigler) focus on privacy as the concealment of (mainly negative) personal information (Posner 1978). Such concealment is assumed to be deliberate and rational: under standard market conditions, the amount of personal information that will be revealed during a transaction merely depends on the trade-off associated with protection of privacy and disclosure of personal information (Stigler 1980) for each party involved (the holder and the potential recipient of data). According to this perspective, individuals can be relied on to rationally seek enough privacy to conceal, and to share, the optimal amount of personal information.

However, while privacy decision making is, no doubt, partly strategic, there are reasons to believe that individuals' preferences for privacy may not be as stable or as internally consistent as the standard economic perspective assumes. The costs of violations of privacy are often amorphous (for example, how bad is it for another person to get a glimpse of one's naked body? What if someone knows what you purchased yesterday on Amazon.com?). And even when the economic costs of such

1. Consider, among others, the Data Security and Breach Notification Act of 2011, S. 1207, 112th Cong. (2011); the Commercial Privacy Bill of Rights Act of 2011, S. 799, 112th Cong. (2011); the Personal Data Privacy and Security Act of 2011, S. 1151, 112th Cong. (2011); the Data Breach Notification Act, S. 1408, 112th Cong. (2011); the Personal Data Protection and Breach Accountability Act of 2011, S. 1535, 112th Cong. (2011); the Secure and Fortify Electronic Data Act of 2011, H.R. 2577, 112th Cong. (2011); and the Cybersecurity Enhancement Act of 2011, H.R. 2096, 112th Cong. (2011).

violations are quantifiable because they lead to some tangible cost, the magnitude, timing, and risk of incurring this cost are often uncertain and difficult to assess (Acquisti 2004). It would therefore be reasonable to conjecture that valuations of privacy will be subject to many of the effects that have come under the heading of preference uncertainty (Slovic 1995). When preferences are uncertain, research has shown, decision making is likely to be influenced by factors that are difficult to justify on normative bases, such as how alternatives are framed (Tversky and Kahneman 1974) or preferences are elicited (Tversky, Slovic, and Kahneman 1990).

We apply theories from behavioral economics and decision research to investigate, and ultimately challenge, the premise that privacy valuations can be precisely estimated. We do so using a field experiment in which we assess the stability of the value that individuals assign to the protection of their personal information. We show, empirically, that privacy valuations are affected not only by endowment (Kahneman and Tversky 1979) but also by the order in which different privacy options are described (Schwarz 1999). In documenting these two effects, we highlight, more generally, that privacy valuations are highly sensitive to nonnormative influences—factors that, in principle, should not affect decision making.

The results of the experiment challenge the robustness of estimates of privacy valuations proposed in the literature and call into question the common conclusion that consumers do not care for privacy: whether they appear to care a lot or a little depends critically on context. The results suggest that discussions about privacy valuations often conflate two different types of transactions that individuals face: transactions in which individuals are offered tangible or intangible benefits in exchange for their personal information and transactions in which individuals are offered protection of their personal information but at some tangible or intangible cost. In our experiment, subjects were five times more likely to reject cash offers for their data if they believed that their privacy would be, by default, protected than if they did not have such a belief. Our findings suggest a vicious (or virtuous) circle of privacy valuations of potential interest to policy makers: those who feel they have less (more) privacy tend to value privacy less (more) and become more (less) likely to accept monetary offers for their data; giving away (protecting) their data, in turn, may make individuals feel they have less (more) privacy—and so on. Such findings highlight the value of insights from

behavioral economic research to properly understand privacy decision making and inform the policy debate surrounding privacy.

2. BACKGROUND AND HYPOTHESES

The empirical literature on privacy valuations is closely connected to the theoretical literature on the economics of information. Economists became interested in studying how agents negotiate privacy trade-offs, and the consequences of their decisions, beginning in the late 1970s with the contributions of Hirshleifer (1980) and Chicago School scholars such as Posner (1978, 1981) and Stigler (1980). Renewed interest in this area arose around the mid-1990s (see, for instance, Varian 1997; Noam 1997; Laudon 1996). In more recent years, formal microeconomic models of privacy trade-offs started appearing (see, for instance, Taylor 2004; Acquisti and Varian 2005; Calzolari and Pavan 2006; Tang, Hu, and Smith 2008; Hann et al. 2008). At the same time, the management, marketing, legal, and information systems literatures also explored the concept of a privacy calculus—such as the anticipation and comparison of benefits, costs, and other consequences associated with the protection of private information (see, for instance, Laufer and Wolfe 1977; Stone and Stone 1990; Culnan and Armstrong 1999; Dinev and Hart 2006).

Implicit in most of the neoclassical economics literature on privacy is the assumption that consumers are rationally informed agents with stable privacy preferences (see, for instance, Posner 1978; Stigler 1980). Most models also assume that privacy is not valued *per se* but for some type of economic benefit it confers. For example, some models focus on consumers' desire to not reveal their personal preferences to a merchant so as to avoid price discrimination in a repeated-purchase scenario (Acquisti and Varian 2005; Taylor 2004). Accordingly, a substantial, and currently active, line of empirical research has attempted to measure individual privacy valuations—an endeavor premised on the assumption that there are, in fact, stable preferences to be measured.

2.1. Estimates of Privacy Valuations

Many empirical efforts in the field of privacy have tried to pinpoint individuals' monetary valuations of privacy. Some of the studies have relied on experiments in which participants faced actual privacy and financial trade-offs, while others have relied on hypothetical surveys (see, for instance, Acquisti 2004). Most of these efforts have focused, either explicitly or implicitly (via the authors' unstated assumptions), on in-

dividuals' willingness to accept payment in exchange for disclosing otherwise private information. Huberman, Adar, and Fine (2005), for example, use a second-price auction to study the amount of money individuals would need to be paid to reveal their weight or height to others. Wathieu and Friedman (2007) show that survey participants were comfortable with an institution's sharing of their personal information if they had been shown the economic benefits of doing so. Cvrcek et al. (2006) find significant differences in the price that European Union citizens would accept to reveal their mobile phone location data that depended on their country of residence. Hui, Teo, and Lee (2007) use a field experiment in Singapore to study the value of various privacy assurance measures and find that privacy statements and monetary incentives could induce individuals to disclose personal information. Chelappa and Sin (2005) also find evidence of a trade-off between consumer valuation for personalization and concerns for privacy. Often, this literature has shown that privacy valuations are low. For example, Tedeschi (2002) reports on a 2002 Jupiter Research study in which 82 percent of online shoppers were willing to give personal data to new shopping sites in exchange for the chance to win \$100. Spiekermann, Grossklags, and Berendt (2001) study individuals' willingness to answer personal questions in order to receive purchase recommendations and discounts and find that individuals who expressed high levels of concern about privacy revealed personal information in exchange for small discounts.

Empirical studies in which consumers are, instead, asked to consider paying (or giving up) money to protect their data are much scarcer. Among those, Rose (2005) finds that although most survey respondents reported that they were concerned about their privacy, only 47 percent of them expressed a willingness to pay any amount to ensure the privacy of their information. Tsai et al. (2011) find that when privacy-relevant information was made salient, participants in an experiment paid moderate premia to purchase both privacy-sensitive and nonsensitive goods from online merchants with better privacy protection; however, when privacy information was not made salient, participants would not pay such premia. The first result was replicated in a more recent European field study by Jentzsch, Preibusch, and Harasser (2012), while the second result was also found in a more recent study by Beresford, Kübler, and Preibusch (2012). Varian, Wallenberg, and Woroeh (2005) and Png (2007) try to estimate the implicit price that U.S. consumers would pay for the protection from telemarketers and find values ranging from a few cents to slightly more than \$30.

In the language of economics, the first set of studies focused on individuals' willingness to accept (WTA): the lowest price a person would be willing to accept to part with a good (protection of personal data) she initially owned. The second set of studies focused on individuals' willingness to pay (WTP): the maximum price a person would be willing to pay to acquire a good (protection of personal data) she did not own. In the privacy literature, these two standpoints are treated as equivalent. However, outside this realm, economic experiments have uncovered a dichotomy between WTP and WTA: WTA tends to be larger than WTP (Hammack and Brown 1974; Kahneman 1986; Knetsch 1989; Kahneman, Knetsch, and Thaler 1990, 1991) for a vast array of both tangible and intangible goods (see, for instance, Dubourg, Jones-Lee, and Loomes 1994). Although various explanations have been proposed for this WTP-WTA gap (Hanemann 1991; Hoehn and Randall 1987), loss aversion—the disproportionate weight that people tend to place on losses relative to gains—is by far the best supported (Kahneman and Tversky 1979; Thaler 1980).²

Applied to privacy, this explanation of the WTA-WTP gap would predict that someone who enjoyed a particular level of privacy but was asked to pay to increase it would be deterred from doing so by the prospect of the loss of money, whereas someone who was asked to sacrifice privacy for a gain in money would also be reluctant to make the change, deterred in this case by the loss of privacy.

Surprisingly, while presenting their results as empirical estimates of individuals' valuations for privacy, none of the empirical studies of privacy valuations have explicitly contrasted individuals' willingness to pay to protect data with their willingness to accept money to reveal the same data. In fact, the very distinction between the two concepts is absent in the literature. For instance, Hann et al. (2007, p. 14) use conjoint analysis to quantify the value individuals ascribe to Web site privacy protection and conclude that “among U.S. subjects, protection against er-

2. An alternative account of the gap between willingness to act and willingness to pay that has been garnering substantial recent support (Weaver and Frederick 2012; Isoni 2011) attributes it to the desire of sellers to not sell at a price below the fair-market price and the desire of buyers to not pay more than the market price. Application of this account to a situation such as privacy, in which there are no established market prices, is difficult. However, Shane Frederick (professor of marketing at Yale School of Management, e-mail to Loewenstein, April 9, 2013) proposed the following conceptually related account: “[S]ellers use high values to signal that their dignity is not for sale, and buyers use low values to signal their refusal to accept the implication that they are entitled to only intermediate levels of privacy.”

rors, improper access, and secondary use of personal information is worth US\$30.49–44.62.” Hann et al. (2007) is a seminal contribution in this area: it offers a first insight, and quantification, of the value individuals assign to online privacy, and in doing so it also stimulates more research in this area. However, conjoint analyses have not distinguished between how much people will pay to protect their data and how much they will accept to give their data away. If it is established that these values differ, then such studies can conclusively determine neither the value of protection against errors nor the true estimate of the value that individuals assign to data. A similar problem exists with the revealed-preferences approach used in other studies. By contrast, in our experiment, one out of two individuals primed to believe that their privacy was, by default, protected rejected cash offers for their data, but few were willing to sacrifice an equivalent amount of cash to prevent the release of the same data. Which of these valuations should be considered the true value of the privacy of our data? Both cannot simultaneously reflect our true preferences.

The distinction between WTP and WTA seems critical in the privacy realm, because real-life, everyday privacy decisions come in both varieties. Analogous to WTP, people are regularly faced with opportunities to pay to prevent personal data from being disclosed—for example, using an anonymous Web-browsing application, such as Tor,³ hides one’s online behavior but incurs the cost of slower downloads; likewise, deleting cookies to shield one’s browsing habits comes at the cost of having to frequently provide registration information across a variety of Web sites. In other situations, analogous to WTA, people are asked to reveal personal information in exchange for some financial benefit—for example, the Internet data company comScore offers its panelists a bundle of products in exchange for monitoring their Internet behavior,⁴ and various loyalty programs offer discounts or awards in exchange for longer and more accurate data trails documenting consumer behavior.

Behavioral decision research tells us that the problem of constructing reliable mappings of consumers’ preferences is not unusual: it applies to a majority of ordinary goods. For such goods, however, markets exist in which the items are bought and sold by consumers and, therefore, objective prices are formed. In the case of privacy, however, consumers by and large do not participate in (and frequently remain unaware of)

3. Tor (<https://www.torproject.org/>).

4. ComScore (<http://www.comscore.com/>).

the daily trades involving their personal data: “infomediaries” such as Choicepoint or credit-reporting agencies such as Experian make a business of buying, aggregating, and selling consumer data (from Social Security numbers to purchasing habits; from financial to medical records) to and from public and private organizations. Only a fraction of those data are made available to, or can be managed by, the consumers who generated them (for instance, redacted credit reports). Of course, consumers do make frequent (almost continuous) decisions involving the protection, or sharing, of their personal information, but these decisions are predominantly bundled into (and therefore both influenced and hidden by) larger economic transactions. For example, the decision to use a pharmacy loyalty card (which creates a record of potentially sensitive purchases at a given store in exchange for a monetary discount on the items purchased) is attached to the completion of pharmacy shopping, which makes it hard to separate consumers’ valuations of privacy from their valuation of discounts and purchased goods. These types of trade-offs are becoming more common, even inescapable: in some cases, consumers can get access to certain goods or services (such as listening to music on Spotify⁵ or commenting on news stories on the *Los Angeles Times*’s Web site) only through a social network that tracks their behavior and links it to their actual identities (Facebook).

The dichotomy between WTP and WTA is just one example of the notion that preference for privacy may be not only context dependent but malleable and uncertain and suggests that ordinary studies investigating privacy valuations may not tell us much about whether, or how much, consumers will actually pay to protect their data. Behavioral economists have highlighted that nonnormative factors often affect valuations and decision making under uncertainty (Slovic 1995). Since many privacy decisions take place under those conditions, researchers have started investigating the impact of cognitive and behavioral biases (on hyperbolic discounting, see Acquisti [2004]; on the illusion of control, see Brandimarte et al. [2010]; on the effect of context on the propensity to disclose, see John et al. [2011]) on privacy decisions and how those decisions deviate from the patterns of behavior predicted by traditional neoclassical economic theory.

In Section 2.2, we formalize how theories from behavioral economics and decision research may apply to privacy and influence both the way individuals value the protection of their personal information and, there-

5. Spotify (<http://www.spotify.com/>).

fore, the extent to which researchers are able to measure those valuations.

2.2. Hypotheses

Consider a consumer with a utility function $u(w, p)$ defined over wealth and privacy. Assume, further, that p^+ represents a situation with greater privacy protection than p^- . For example, p^- might represent a purchase completed via an ordinary credit card, while p^+ could represent the same purchase made with an anonymous payment method (from cash to more sophisticated technologies such as those described in Chaum [1983]). For individuals who begin in the position $u(w, p^+)$, the smallest amount they should be willing to accept to shift to p^- is given by the equation $u(w + \text{WTA}, p^-) = u(w, p^+)$. Likewise, for individuals who begin in situation p^- , the most they should be willing to pay to shift to a situation characterized by p^+ is $u(w - \text{WTP}, p^+) = u(w, p^-)$. The implication of these equations is that WTA will not necessarily be identical to WTP, and, in particular, if privacy is a normal good that becomes valued more as one becomes wealthier, it is possible that WTA is greater than WTP, although one would expect the difference to be trivial given almost any plausible form of the utility function (Willig 1976). Nevertheless, as the equations show, the existence of a discrepancy between WTA and WTP cannot in and of itself be viewed as a violation of standard economic theory.

Suppose, however, that the individuals in the two situations are faced with binary trade-offs between privacy and money, with monetary transfers creating two possible final levels of wealth: w^+ and w^- , with $w^+ > w^-$. In WTA mode, the consumer faces a choice between an initial position of w^- and p^+ and the choice of obtaining money in exchange for reduced privacy, which leads to w^+ and p^- . In WTP mode, the consumer faces a choice between an initial position of w^+ and p^- and the choice of paying to gain greater privacy, which leads to w^- and p^+ . Whether the first consumer will choose to accept the payment will depend on whether $u(w^-, p^+) < u(w^+, p^-)$. Whether the second consumer will choose to pay the fee will depend on whether $u(w^+, p^-) > u(w^-, p^+)$. Clearly, these conditions are precisely the same. Thus, standard economic theory predicts that people will make identical choices in these two situations, regardless of whether they are framed in terms of WTA (a loss of privacy and gain of money) or WTP (a gain of privacy and loss of money).

To provide a clean test of nonnormative differences between WTP

and WTA, therefore, we elicited privacy preferences through binary choices that were identical from the perspective of final outcomes but differed in initial endowments. Such binary choices are characteristic of many real-world situations. Consumers are rarely asked how much they would be willing to pay (need to be paid) for (to avoid) some change in privacy. Instead, they are typically given binary choices, including take-it-or-leave-it options. For example, choosing to use a grocery loyalty card (which tracks individual purchases but offers a discount the consumers cannot negotiate) or not, choosing to use Pretty Good Privacy encryption (which protects e-mail content but is harder—and therefore costlier—to use) or not, and so forth. A rational consumer conforming to the dictates of standard economics would display similar preferences regardless of whether a choice was framed in terms of WTA or WTP. However, if consumers were affected by a sense of endowment in the privacy of their data, their preferences facing those two choices would be different. Accordingly, we make the following hypothesis.

Hypothesis 1: Willingness to Pay and Willingness to Accept Payment for Privacy. The fraction of consumers who will reject an offer to obtain money in exchange for reduced privacy (WTA) is larger than the fraction of consumers who will accept an economically equivalent offer to pay money in exchange for increased privacy (WTP).

If this hypothesis is correct, it would imply the possibility that $u(w^-, p^+) > u(w^+, p^-)$ while also, simultaneously, that $u(w^+, p^-) > u(w^-, p^+)$, simply depending on how the question is framed. This would suggest that the minimum price a consumer will be willing to accept to allow personal data to be revealed may be higher than the maximum price she or he will be willing to pay to avoid having that data revealed—in other words, consumers may value their personal information more when they are endowed with it (namely, with its protection) and are asked to reveal it than when they begin without protection and are given the opportunity to pay to obtain it. This would also suggest, more broadly, that privacy preferences, while not arbitrary, are malleable to nonnormative factors and can be internally inconsistent, in that the same cash-for-data offer may be accepted or rejected for nonnormative reasons.

Another aspect of privacy valuations worth considering is that if privacy costs and benefits are difficult to estimate with any precision, individuals may form their valuations of privacy on the basis of contextual cues with little normative justification. Consider, in particular,

the fact that consumers' decisions are often affected by the order in which offers are presented (Brookshire et al. 1981; Schwarz 1999; in related work, Johnson, Bellman, and Lohse [2002] studied default effects in privacy decision making). Applied to privacy, this anomaly would suggest that consumers' privacy valuations could depend on the order in which they are asked to reveal privacy-sensitive information. Hence, we predicted that presenting a privacy-enhanced option prior to one that is relatively less protective of privacy may be interpreted as a signal that the former is inherently more valuable.

Hypothesis 2: Order Effects in Privacy Valuations. Faced with the choice between offers with different monetary values and privacy features, the fraction of consumers who will choose a privacy-enhanced offer is larger when that offer is presented before its (less privacy protective) alternative.

3. FIELD EXPERIMENT

We used a field experiment to test our hypotheses. Subjects were asked to choose between gift cards that varied with respect to their privacy features and monetary values. We focused on informational privacy—concerns over the treatment of one's purchase data (Tsai et al. 2011). We investigated subjects' willingness to keep versus exchange gift cards as a function of (1) their initial endowment and (2) the order in which choices were presented. The experiment tested hypotheses 1 and 2 in the field with real gift cards. Our scenario focuses on U.S. consumers and, in particular, their willingness to disclose purchase transaction data with researchers.

Subjects were offered Visa gift cards that could be used to purchase goods from any online or offline store where debit cards are accepted. Shopping mall patrons were stopped by research assistants (blind to the hypotheses of the study) and offered gift cards in exchange for participating in a survey. In reality, the survey was a decoy intended to create a credible explanation for (and detract attention from) the gift cards that subjects were given as a reward. Across all conditions, subjects had to choose between the same two alternatives: a \$10 anonymous card and a \$12 identified card. For the former card, subjects were told that their "name will not be linked to the transactions completed with this card." For the \$12 identified card, they were told that their "name will be linked to the transactions completed with this card." The framing of this choice differed between experimental conditions.

The study was a five-condition between-subjects design. In two endowed conditions, subjects were endowed with either the \$10 anonymous card or the \$12 identified card before being offered the option to swap one card for the other. These conditions were used to test whether, and how significantly, the endowment effect played a role in privacy valuations. In two choice conditions, subjects were not endowed with a particular card before choosing but were simply asked to choose between either a \$10 or \$12 gift card or a \$12 or \$10 gift card (in one condition, the anonymous \$10 card was described first; in the other, it was described second). The choice conditions allowed us to test the role of order effects in privacy valuations but were also included to situate the impact of the WTA and WTP conditions relative to more neutral conditions that did not incorporate a status quo. Finally, we included one rationality check control condition, in which the choice was between a \$10 identified card and a \$12 anonymous card. In this condition, the latter card was both more valuable and more privacy preserving than the \$10 card and thus a clearly dominant choice. This condition was included to ensure that people understood and paid attention to the task. We summarize the four main conditions:

1. \$10 endowed: Keep the anonymous \$10 card or exchange it for an identified \$12 card.
2. \$12 endowed: Keep the identified \$12 card or exchange it for an anonymous \$10 card.
3. \$10 choice: Choose between an anonymous \$10 card and an identified \$12 card.
4. \$12 choice: Choose between an identified \$12 card and an anonymous \$10 card.

All subjects in these four conditions, regardless of the condition to which they had been randomly assigned, faced the same alternatives: a \$10 anonymous card or a \$12 identified card. There was no deception: subjects were asked to choose between gift cards preloaded with money that could be spent at real offline and online merchants. Although the purchases made on both types of cards could be tracked, only purchases made on identified cards could be linked to the person's name.

However, the gift card endowment manipulation generated a different framing of the card choice: for those in the \$10 endowed condition, the question was framed as an implicit choice to sell one's future purchase data to the researchers for \$2; for those in the \$12 endowed condition, the question was framed as an implicit choice to pay \$2 to avoid having

one's future purchase data made available to the researchers.⁶ Since subjects across those conditions faced exactly the same two alternatives, the percentages of people choosing the anonymous card over the identified one should remain the same, regardless of the framing. If those percentages differed, this would provide evidence of a WTP/WTB dichotomy and/or order effects.⁷

Per the guidelines introduced by Simmons, Nelson, and Simonsohn (2011), we report how we determined our sample size, all manipulations, and all measures in the study. No data were excluded.

3.1. Procedure

The experimental procedure is summarized here (complete details are available in the online Appendix). The experiment took place at a shopping mall in the Pittsburgh area. We collected as much data as we could over 3 weekend days. Female research assistants stood at the entrance of two women's clothing stores and approached female shoppers as they entered and asked them to complete a brief survey designed to assess people's attitudes toward spending money. Interested shoppers were given a coupon valid for a gift card, redeemable on exiting the store, for completion of a short survey. After completing the survey and on exiting the store, each subject gave her coupon to the experimenter, who then asked the subject (regardless of condition) to print her name at the top of a receipt for the gift card. The experimenter then called the subject by her name, informing her that the coupon was valid for a gift card. Subjects were addressed by their names to increase the salience of the name-identification feature of the identified gift cards. Next, the experimenter gave the subject a sheet of paper, noting that it outlined the features of the card. Experimenters were trained to avoid words such as "tracked" and "privacy" that may have alerted subjects to the purpose of the study.

6. We designed the experiment to focus on future transaction data (that is, a piece of personal information that did not yet exist at the time subjects had to choose between the cards) to avoid potential confounds associated with potential previous disclosures of existing personal data.

7. Naturally, if a subject's valuation of her personal data were, for instance, \$.50, it would be rational for her to switch to a trackable card for \$12 (from a \$10 untrackable card) in one condition and to accept a \$12 trackable card in a different condition. But since subjects with various heterogeneous privacy valuations were randomly assigned to the conditions, we can expect *ex ante* privacy valuations to be also similarly distributed. In such a case, the proportion of people who choose the trackable card over the untrackable card should also remain the same across conditions.

Until this point, subjects across the five conditions had been exposed to the same experience, and all had provided the same amount of personally identifying information to the researchers. Thereafter, subjects in the endowed conditions were given a sheet that described the features of the card with which they were to be endowed. The subject then selected a card from the appropriate bin, be it the \$10 or \$12 gift card bin. Next, the experimenter gave the subject a second sheet of paper describing the privacy features of the other card. The subject was then asked whether she would like to exchange her \$10 anonymous (\$12 identified) card for the \$12 identified (\$10 anonymous) card. All subjects—even those who chose to keep the card with which they had been endowed—were required to check a box on the receipt that coincided with their final gift card choice. As subjects decided whether to exchange their gift cards, the two bins of gift cards were in plain sight and within arms' reach. Therefore, any costs associated with switching cards were trivial: a subject who wanted to switch cards simply dropped her initial card into the appropriate bin and selected a new card from the other bin. To the extent that switching costs existed, they were kept constant across conditions.

In the choice conditions, subjects were presented with only one description sheet that listed and described both cards, one after the other, with the order of description presentation manipulated between subjects. Each subject then indicated which card she would like and selected a card from the appropriate bin. Each subject was then asked to provide her e-mail address.

All subjects had the same amount of time to reflect on how to use their cards in the future. In particular, all subjects, regardless of their experimental condition, could have mentally compared choosing the trackable card to purchase nonsensitive items versus choosing the anonymous card to purchase more privacy-sensitive items.

3.2. Results

A total of 349 female subjects participated in the study (mean age = 35; median age = 35; 83.6 percent Caucasian; all not significant between conditions). On exiting the store, the majority (92.3 percent) of subjects returned to the experimenter to redeem the gift card coupon.⁸

8. The results presented in Section 3.2, which are based on a field experiment, are also consistent with those of additional survey experiments, which are available from the authors on request.

Figure 1. Gift card selection

3.2.1. Gift Card Choice. Virtually everyone in the rationality check control condition (95.7 percent) selected the \$12 anonymous card, which suggests that subjects understood and took the task seriously.⁹ This condition is excluded from the rest of the analysis.

The proportion of people choosing the \$10 anonymous card was highest when subjects had been endowed with it (52.1 percent). Next highest was for the choice condition in which the \$10 card was listed first (42.2 percent), followed by the choice condition in which the \$10 card was listed second (26.7 percent). Lowest (9.7 percent) was for those endowed with the \$12 identified card (see Figure 1).¹⁰

Subjects in the endowed conditions displayed a tendency to keep the card with which they had been endowed, which supports previous results on the power of default settings on privacy decision making (Johnson, Bellman, and Lohse 2002). However, and more interestingly, while 90.3 percent of subjects in the \$12 endowed condition kept the \$12 card,

9. This result shows preference for both money and anonymity, which suggests that subjects preferred to keep their transaction data hidden from the researchers.

10. Figure 1 shows that in the \$10 endowed condition, 37 subjects chose the \$10 card and 34 subjects chose the \$12 card. In the \$10 choice condition, the number of subjects choosing each card was, respectively, 35 and 48. In the \$12 choice condition, the numbers were 16 and 44. In the \$12 endowed condition, the numbers were 6 and 56.

Table 1. Probit Regressions by Condition

	Endowed	Choice
Constant	2.4379** (.4880)	1.1130** (.3608)
Age	-.0304** (.0104)	-.0102 (.0082)
\$10Card	-1.4400** (.2917)	-.6210* (.2417)
N	123	128
Pr > $\chi^2(3)$	.0000	.0180
Pseudo- R^2	.23	.05

Note. The dependent variable represents the card selection (\$10 anonymous card = zero, \$12 identified card = one). Standard errors are in parentheses.

* $P < .10$.

** $P < .01$.

only 52.1 percent of those in the \$10 endowed condition kept the \$10 card. In other words, significantly more subjects in the \$12 endowed condition kept their card than those in the \$10 endowed condition ($\chi^2(1) = 27.24$; $p < .001$). The results of the two choice conditions—differing only in the order in which the cards were described—are marginally significantly different from each other ($\chi^2(1) = 3.64$; $p = .056$): subjects seemed more likely to choose the card that was described first. In particular, when the \$12 identified card was listed first, 73.3 percent of subjects chose it, whereas when it was listed after the description of the \$10 anonymous card, only 57.8 percent of subjects chose it.

Table 1 presents the results of two logistic regressions in which we regressed age and dummy variables representing the experimental conditions over a dichotomous dependent variable representing the selection of the traditional \$12 gift card (one) over the privacy-enhanced \$10 gift card (zero).¹¹ We ran one regression for the two endowed conditions and one for the two choice conditions. We used a dummy variable (\$10Card) to control for which card the subject was endowed with (or presented first): the \$10 card (one) or the \$12 card (zero). The result of both models is significant. In the endowed conditions, the coefficient on \$10Card is strongly significant and negative ($p < .001$). This result strongly supports hypothesis 1. In the choice conditions, the coefficient

11. Sheehan (1999, 2002) highlights age and gender differences in privacy concerns. Therefore, we control for age in the regression analysis. We did not use a dummy for gender since, as noted, the experiment focused on a female population.

on \$10Card is negative and weakly significant ($p = .1$), which provides modest support for hypothesis 2 and also indicates that order effects are weaker than endowment effects.

3.2.2. Card Usage. We tracked the stores at which subjects used their gift cards to make purchases (although we could not ascertain what products they purchased). One month after the study, the majority of subjects (87.7 percent) had used their cards. Subjects who had chosen the more valuable card were slightly more likely to have used it (90.7 percent of those with \$12 cards versus 81.8 percent of those with \$10 cards; Pearson $\chi^2(1) = 4.25$; $p = .039$). There were no significant differences in the propensity to use the card with respect to the initial conditions of assignment: whether the subject had been initially endowed with, or had to initially choose, a card (Pearson $\chi^2(1) = .16$; $p = .688$) or whether the subject had been initially assigned an anonymous or identified card (Pearson $\chi^2(1) = 1.28$; $p = .258$).

We investigate whether subjects used their cards at different types of stores depending on card identifiability. Stores were classified as potentially privacy sensitive (for example, lingerie stores such as Victoria's Secret) or not (cafes, convenience stores, supermarkets). We find suggestive, although by no means definitive, evidence of differences in store patronage depending on card identifiability. For instance, all of the eight purchases recorded at Victoria's Secret were completed with the more valuable but less privacy protected card. Future research could test whether having the option to pay a premium for privacy-protecting cards is associated with, or even induces, more privacy-sensitive purchases.

3.2.3. Explanations for Decisions. In the exit questionnaire, we asked subjects to explain why they chose one card over the other. Explanations provided by subjects who chose the \$10 card often referenced privacy concerns and in particular a resistance to being tracked: "Didn't want to give name," "Didn't want to be linked," "[Wanted] privacy," "Didn't want to disclose my information," and "Would rather it be anonymous." Only one subject referred to actual risks by noting that "[the \$10 card] seemed to be safer." In contrast, subjects who chose the \$12 card mostly explained their choice using variations of "More money to spend!" or "Because it was more money!" or even referred in particular to not fearing being tracked: "I don't mind if people know what I buy," "It doesn't bother me if you know where I spend it," and "I don't mind if you know where I spend my money."

4. DISCUSSION

Shoppers implicitly assigned dramatically different values to the privacy of their data depending on the framing of the choice between gift cards with different privacy features. The number of subjects willing to reject cash offers for their data was both significant in absolute terms and much larger in relative terms when they felt that their data would be, by default, protected (\$10 endowed condition) than when they believed that their data would be, by default, revealed (\$12 endowed condition). Implicit valuations of privacy were also affected by the order in which the gift cards were described—the fraction of consumers who chose the privacy-enhanced card was larger when that card was presented before its (less privacy protective) alternative.

More than half of subjects who had been endowed with an anonymous \$10 card rejected an offer of \$2 to reveal their future purchase data (that is, an increase of 20 percent of their initial endowment): these subjects decided that \$2 was not enough to give away their privacy, even though they could have planned to use a trackable card in the future for non-privacy-sensitive transactions. The WTA of these individuals was therefore larger than (or at best equal to) \$2. By contrast, fewer than 10 percent of subjects endowed with the identified \$12 card chose to give up \$2 (a 17 percent decrease in their initial endowment) to protect future purchase data. The overwhelming majority of these subjects refused to pay \$2 to protect their future purchase data—they decided that \$2 was too much to protect their privacy. Subjects were five times more likely to choose privacy in one condition over the other, even though all subjects faced exactly the same choice. Although our experiment was conducted in a U.S. setting and constrained to a sample of female shoppers, it is consistent with a battery of additional survey experiments with alternative populations.

Making some simplifying assumptions, we can compare the privacy WTA:WTP ratio to similar ratios estimated in the literature for other private goods. Let us assume that, *ex ante*, subjective privacy valuations were clustered at \$0 for those who opted to share information and \$2 for those who did not (note that choosing values higher than \$2 would merely increase estimated differences between conditions). Then, the *ex post* mean valuation in the \$10 endowed condition could be calculated at roughly \$1.04 ($[(.52 \times \$2) + (.48 \times \$0)]$), and that in the \$12 endowed condition could be calculated at roughly \$.19. These results represent a WTA:WTP ratio of 5.47—markedly larger than the average ratio ob-

servable for ordinary private goods (which Horowitz and McConnell [2002] report as 2.92).

Such a gap between privacy WTP and WTA is notable because, while ordinary consumer goods (whose valuations can also be affected by the endowment effect) are directly traded in markets where objective prices are formed, privacy transactions are most often bundled with other primary transactions, which makes the estimation of privacy valuations for the benefits of public policy and decision making even more challenging. In everyday life, individuals face privacy trade-offs of two types that often get conflated in policy and empirical debates about the value of privacy: transactions in which individuals are offered tangible or intangible benefits in exchange for their personal information and transactions in which individuals are offered protection of their personal information but at some tangible or intangible costs.

At an empirical level, our findings should caution against the uncritical use of privacy valuations that have used single methods—for example, only WTP or only WTA. Such often surprisingly precise valuations should be interpreted with extreme caution: failing to differentiate between how much an individual would pay versus accept for private data conceals the reality of how malleable and mutable these valuations can be. The answers to questions such as What is privacy worth? and Do people really care for privacy? depend not just on whom, but how, you ask.

From a theoretical standpoint, we show that the assumption that privacy valuations are independent of endowment is empirically questionable. Since economic models are used to influence and direct public policy initiatives, our empirical results may carry a practical lesson to guide our efforts as modelers: our models should account for the fact that estimated valuations of privacy depend on the direction of the cash-for-privacy exchange: they are larger when individuals consider trading personal data for money and smaller when people pay money for privacy.

Finally, and perhaps most important, from a policy perspective, this research raises the issue of individuals' abilities to optimally navigate issues of privacy. From choosing whether to join a grocery loyalty program to sharing sensitive information (such as one's Social Security number) with a merchant, individuals make frequent privacy-relevant decisions, and this research suggests that they do so inconsistently. In the debate surrounding privacy in the United States, great attention has been paid to notice-and-consent solutions that provide increased transparency and control to individuals about what happens to their personal infor-

mation.¹² Our findings raise the question of whether notice-and-consent solutions (or similar self-regulatory approaches) may be sufficient to guarantee consumers' privacy protection. Very often, in online settings, users' decisions are affected by defaults chosen by the providers of Internet services or embodied in the architecture of Internet technologies, which can create either a WTP or a WTA transaction for consumers' data. For instance, certain fields of personal data on popular social networking sites (such as Facebook or Google+) are by default set to be private, while others are set to be public. The most popular Internet browsers by default leak users' information (such as Internet protocol [IP] addresses, operating systems, referral pages, and so forth) to the servers of the sites they visit; and using a search engine such as Google for a query or a music service such as Spotify to listen to music automatically provides personal information (such as one's searches, browsing habits, or music preferences, which can be linked to an IP address and, increasingly often, to an actual identity) to those services. To avoid revealing personal data, a user should seek, install, and learn to use alternative tools (such as Tor), albeit sometimes at a cost (for instance, slower Internet browsing).

The importance of privacy defaults is perhaps nowhere more apparent than in the current debate over the so-called Do Not Track list (see Tene and Polonetsky 2012). Representatives from the data industry, by and large, do not want to support solutions in which browsers by default treat users as not trackable by behavioral advertisers, whereas consumer advocates (as well as a few industry players) are supportive of default browser settings that do not allow tracking. The finding that endowment effects powerfully influence individual privacy valuations may help to justify the introduction of policy interventions that protect people from their own suboptimal decisions.

Individuals' decisions about their data are sometimes taken as representing true and final preferences toward protection or revelation of personal data and therefore become an instrument for the assignment of societal resources to privacy issues. For example, the observation that individuals give away their personal information for small rewards has permeated the policy debate and has been used to argue against privacy regulation (for example, Rubin and Lenard 2002) on the grounds that

12. Consider, for instance, Senator John Kerry's proposal for a privacy bill (Commercial Privacy Bill of Rights Act of 2011, S. 799, 112th Cong. [2011]) and a privacy report by the U.S. Federal Trade Commission (2010).

if consumers wanted more privacy they would ask for it and take advantage of opportunities to protect it. However, as we have shown, revealed-preferences arguments should not, alone, justify the uncritical conclusion that privacy-conscious consumers will never pay for privacy. If individual privacy decisions are so malleable to endowment and order effects, such arguments lose their normative standing.

REFERENCES

- Acquisti, Alessandro. 2004. Privacy in Electronic Commerce and the Economics of Immediate Gratification. Pp. 21–29 in *Proceedings of the Fifth ACM Conference on Electronic Commerce*, edited by Jack Breese, Joan Feigenbaum, and Margo Seltzer. New York: Association for Computing Machinery.
- Acquisti, Alessandro, and Hal Varian. 2005. Conditioning Prices on Purchase History. *Marketing Science* 24:1–15.
- Barnes, Robert. 2012. Supreme Court Restricts Police GPS Tracking. *Washington Post*, January 24.
- Beresford, Alastair R., Dorothea Kübler, and Sören Preibusch. 2012. Unwillingness to Pay for Privacy: A Field Experiment. *Economics Letters* 117:25–27.
- Brandimarte, Laura, Alessandro Acquisti, and George Loewenstein. 2010. Privacy Concerns and Information Disclosure: An Illusion of Control Hypothesis. Paper presented at the Conference on Information Systems and Technology (CIST), Austin, Tex., November 6–7.
- Brookshire, David S., Ralph C. d'Arge, William D. Schulze, and Mark A. Thayer. 1981. Experiments in Valuing Public Goods. Pp. 123–72 in *Advances in Applied Microeconomics: Volume 1*, edited by V. Kerry Smith. Greenwich, Conn.: JAI Press.
- Calzolari, Giacomo, and Alessandro Pavan. 2006. On the Optimality of Privacy in Sequential Contracting. *Journal of Economic Theory* 130:168–204.
- Chaum, David 1983. Blind Signatures for Untraceable Payments. Pp. 199–203 in *Advances in Cryptology: Proceedings of Crypto '82*, edited by David Chaum, Ronald L. Rivest, and Alan T. Sherman. New York: Springer Verlag.
- Chellapa, Ramnath K., and Raymond G. Sin. 2005. Personalization versus Privacy: An Empirical Examination of the Online Consumers' Dilemma. *Information Technology and Management* 6:181–202.
- Culnan, Mary J., and Pamela K. Armstrong. 1999. Information Privacy Concerns, Procedural Fairness, and Impersonal Trust: An Empirical Investigation. *Organization Science* 10:104–15.
- Cvrcek, Dan, Marek Kumpost, Vashek Matyas, and George Danezis. 2006. A Study on the Value of Location Privacy. Pp. 109–18 in *Proceedings of the 5th ACM Workshop on Privacy in the Electronic Society (WPES '06)*, edited

- by Ari Juels and Marianne Winslett. New York: Association for Computing Machinery.
- Dinev, Tamara, and Paul Hart. 2006. An Extended Privacy Calculus Model for E-Commerce Transactions. *Information Systems Research* 17:61–80.
- Dubourg, W. Richard, Michael W. Jones-Lee, and Graham Loomes. 1994. Imprecise Preferences and the WTP-WTA Disparity. *Journal of Risk and Uncertainty* 9:115–33.
- Gonsalves, Antone. 2010. Facebook CEO: Less Privacy Is Social Norm. *InformationWeek*, January 12.
- Hammack, Judd, and Gardner Mallard Brown. 1974. *Waterfowl and Wetlands: Toward Bioeconomic Analysis*. Baltimore: Johns Hopkins University Press.
- Hanemann, W. Michael. 1991. Willingness to Pay and Willingness to Accept: How Much Can They Differ? *American Economic Review* 81:635–47.
- Hann, Il-Horn, Kai-Lung Hui, Sang-Yong Tom Lee, and Ivan P. L. Png. 2007. Overcoming Information Privacy Concerns: An Information Processing Theory Approach. *Journal of Management Information Systems* 24:13–42.
- . 2008. Consumer Privacy and Marketing Avoidance: A Static Model. *Management Science* 54:1094–1103.
- Hirshleifer, Jack. 1980. Privacy: Its Origin, Function, and Future. *Journal of Legal Studies* 9:649–66.
- Hoehn, John P., and Alan Randall. 1987. A Satisfactory Benefit Cost Indicator from Contingent Valuation. *Journal of Environment, Economics, and Management* 14:226–47.
- Horowitz, John K., and Kenneth E. McConnell. 2002. A Review of WTA/WTP Studies. *Journal of Environmental Economics and Management* 44:426–47.
- Huberman, Bernardo A., Eytan Adar, and Leslie Fine. 2005. Valuating Privacy. *IEEE Security and Privacy* 3:22–25.
- Hui, Kai-Lung, H-H. Teo, and Sang-Yong Lee. 2007. The Value of Privacy Assurance: An Exploratory Field Experiment. *MIS Quarterly* 31:19–33.
- Isoni, Andrea. 2011. The Willingness-to-Accept/Willingness-to-Pay Disparity in Repeated Markets: Loss Aversion or “Bad-Deal” Aversion? *Theory and Decision* 71:409–30.
- Jentzsch, Nicola, Sören Preibusch, and Andreas Harasser. 2012. *Study on Monetising Privacy: An Economic Model for Pricing Personal Information*. Report for the European Network and Information Security Agency. Heraklion: European Network and Information Security Agency.
- John, Leslie K., Alessandro Acquisti, and George Loewenstein. 2011. Strangers on a Plane: Context-Dependent Willingness to Divulge Sensitive Information. *Journal of Consumer Research* 37:858–73.
- Johnson, Eric J., Steven Bellman, and Gerald Lohse. 2002. Defaults, Framing, and Privacy: Why Opting in—Opting out. *Marketing Letters* 13:5–15.
- Kahneman, Daniel. 1986. Comments on the Contingent Valuation Method. Pp. 185–93 in *Valuing Environmental Goods: An Assessment of the Contingent*

- Valuation Method*, edited by Ronald G. Cummings, David S. Brookshire, and William D. Schulze. Totowa, N.J.: Rowman & Allanheld.
- Kahneman, Daniel, Jack L. Knetsch, and Richard H. Thaler. 1990. Experimental Tests of the Endowment Effect and the Coase Theorem. *Journal of Political Economy* 98:1325–48.
- . 1991. Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias. *Journal of Economic Perspectives* 5:193–206.
- Kahneman, Daniel and Amos Tversky. 1979. Prospect Theory: An Analysis of Decision under Risk. *Econometrica* 47:263–92.
- Kelsey, Joel, and Michael McCauley. 2008. Consumer Reports Poll: Americans Extremely Concerned about Internet Privacy. Consumersunion.org, September 25. <http://consumersunion.org/news/poll-consumers-concerned-about-internet-privacy>.
- Knetsch, Jack L. 1989. The Endowment Effect and Evidence of Nonreversible Indifference Curves. *American Economic Review* 79:1277–84.
- Laudon, Kenneth C. 1996. Markets and Privacy. *Communications of the ACM* 39:92–104.
- Laufer, Robert S., and Maxine Wolfe. 1977. Privacy as a Concept and a Social Issue: A Multidimensional Developmental Theory. *Journal of Social Issues* 33: 22–42.
- Lenard, Thomas M., and Paul H. Rubin. 2010. In Defense of Data: Information and the Costs of Privacy. *Policy and Internet* 2:149–83.
- Mulligan, Deirde K., and Janlori Goldman. 1997. The Limits and the Necessity of Self-Regulation: The Case for Both. Chapter 1G in *Privacy and Self-Regulation in the Information Age*. Washington, D.C.: U.S. Department of Commerce, National Telecommunications and Information Administration.
- Noam, Eli M. 1997. Privacy and Self-Regulation: Markets for Electronic Privacy. Chapter 1B in *Privacy and Self-Regulation in the Information Age*. Washington, D.C.: U.S. Department of Commerce, National Telecommunications and Information Administration.
- Png, Ivan P. L. 2007. On the Value of Privacy from Telemarketing: Evidence from the “Do Not Call” Registry. Working paper. National University of Singapore, School of Business.
- Posner, Richard A. 1978. The Right of Privacy. *Georgia Law Review* 12:393–422.
- . 1981. The Economics of Privacy. *American Economic Review* 71:405–9.
- Romanosky, Sasha, and Alessandro Acquisti. 2009. Privacy Costs and Personal Data Protection: Economic and Legal Perspectives. *Berkeley Technology Law Journal* 24:1061–1102.
- Rose, Ellen A. 2005. Data Users versus Data Subjects: Are Consumers Willing to Pay for Property Rights to Personal Information? *Proceedings of the 38th Annual Hawaii International Conference on System Sciences*. Washington, D.C.: IEEE Computer Society Press. doi:10.1109/HICSS.2005.184.

- Rubin, Paul H., and Thomas M. Lenard. 2002. *Privacy and the Commercial Use of Personal Information*. Washington, D.C.: Progress and Freedom Foundation.
- Schwarz, Norbert. 1999. Self-Reports: How the Questions Shape the Answers. *American Psychologist* 54:93–105.
- Sheehan, Kim Bartel. 1999. An Investigation of Gender Differences in On-line Privacy Concerns and Resultant Behaviors. *Journal of Interactive Marketing* 13:24–38.
- . 2002. Toward a Typology of Internet Users and Online Privacy Concerns. *Information Society* 18:21–32.
- Simmons, Joseph P., Leif D. Nelson, and Uri Simonsohn. 2011. False-Positive Psychology: Undisclosed Flexibility in Data Collection and Analysis Allows Presenting Anything as Significant. *Psychological Science* 22:1–8.
- Slovic, Paul. 1995. The Construction of Preference. *American Psychologist* 50: 364–71.
- Solove, Daniel J. 2004. *The Digital Person: Technology and Privacy in the Information Age*. New York: New York University Press.
- Spiekermann, Sarah, Jen Grossklags, and Bettina Berendt. 2001. E-Privacy in Second Generation E-Commerce: Privacy Preferences versus Actual Behavior. Pp. 38–47 in *Proceedings of the Third ACM Conference on Electronic Commerce*, edited by Michael P. Wellman and Yoav Shoham. New York: Association for Computing Machinery.
- Stigler, George J. 1980. An Introduction to Privacy in Economics and Politics. *Journal of Legal Studies* 9:623–44.
- Stone, Eugene F., and Dianna L. Stone. 1990. Privacy in Organizations: Theoretical Issues, Research Findings, and Protection Mechanisms. Pp. 349–411 in vol. 8 of *Research in Personnel and Human Resources Management*, edited by Gerald R. Ferris and Kendrith M. Rowland. Greenwich: JAI Press.
- Tang, Zhulei, Yu Jeffrey Hu, and Michael D. Smith. 2008. Gaining Trust through Online Privacy Protection: Self-Regulation, Mandatory Standards, or Caveat Emptor. *Journal of Management Information Systems* 24:153–73.
- Taylor, Curtis R. 2004. Consumer Privacy and the Market for Customer Information. *Rand Journal of Economics* 35:631–50.
- Tedeschi, Bob. 2002. Everybody Talks about Online Privacy, but Few Do Anything about It. *New York Times*, June 3.
- Tene, Omer, and Jules Polonetsky. 2012. To Track or “Do Not Track”: Advancing Transparency and Individual Control in Online Behavioral Advertising. *Minnesota Journal of Law, Science, and Technology* 13:281–357.
- Thaler, Richard H. 1980. Toward a Positive Theory of Consumer Choice. *Journal of Economic Behavior and Organization* 1:39–60.
- Tsai, Janice Y., S. Egelman, L. Cranor, and Alessandro Acquisti. 2011. The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study. *Information Systems Research* 22:254–68.

- Tversky, Amos, and Daniel Kahneman. 1974. The Framing of Decisions and the Psychology of Choice. *Science* 211:453–58.
- Tversky, Amos, Paul Slovic, and Daniel Kahneman. 1990. The Causes of Preference Reversal. *American Economic Review* 80:204–17.
- U.S. Department of Commerce. 2010. *Commercial Data Privacy and Innovation in the Internet Economy: A Dynamic Policy Framework*. Internet Policy Task Force Green Paper. Washington, D.C.: U.S. Department of Commerce.
- U.S. Federal Trade Commission. 2010. *Protecting Consumer Privacy in an Era of Rapid Change: A Proposed Framework for Businesses and Policymakers*. Washington, D.C.: U.S. Federal Trade Commission. <http://www.ftc.gov/os/2010/12/101201privacyreport.pdf>.
- Varian, Hal R. 1997. Economic Aspects of Personal Privacy. Chapter 1C in *Privacy and Self-Regulation in the Information Age*. Washington, D.C.: U.S. Department of Commerce, National Telecommunications and Information Administration.
- Varian, Hal R., Fredrik Wallenberg, and Glenn Woroch. 2005. The Demographics of the Do-Not-Call List. *IEEE Security and Privacy* 3:34–39.
- Wathieu, Luc, and Allan Friedman. 2007. An Empirical Approach to Understanding Privacy Valuation. Working Paper No. 07-075. Harvard Business School, Cambridge, Mass.
- Weaver, Ray, and Shane Frederick. 2012. A Reference Price Theory of the Endowment Effect. *Journal of Marketing Research* 49:696–707.
- Willig, Robert D. 1976. Consumer's Surplus without Apology. *American Economic Review* 66:589–97.