

What is privacy worth?

Alessandro Acquisti

(Carnegie Mellon University; acquisti@andrew.cmu.edu)

Leslie John

(Harvard Business School; ljohn@hbs.edu)

George Loewenstein

(Carnegie Mellon University; gl20@andrew.cmu.edu)

What is privacy worth?

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 investigate individual privacy valuations in a series of field and online experiments informed by behavioral economics and decision research, and find evidence of order and endowment effects, and non-normal distributions of privacy valuations. Individuals assign markedly different values to the privacy of their data depending on a) how much money they would accept to disclose otherwise private information, or how much they would pay to protect otherwise public information; and b) the order in which they consider different offers for that data. The gap between such values is high compared to that observed in comparable studies of other private goods. In a series of additional experiments, we test robustness and boundary conditions of our main effects and whether they apply to different types of privacy concerns. The results paint a nuanced picture of privacy valuations and highlight their sensitivity to contextual, non-normative effects.

1. INTRODUCTION¹

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, due in part to the emergence of new technologies, such as GPS tracking and social networking over the Internet. In a recent case described in the Washington Post (Barnes [2012]) 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 drug dealer based on the argument that monitoring the movements of his Jeep by affixing a GPS device to his Jeep for 28 days violated his Fourth Amendment rights. As has often been pointed out, the 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 goal 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 against other desirable goals. For example, the Health Insurance Portability and Accountability Act of 1996 (HIPAA) gave patients greater privacy protections than they had before, but at the cost of increased administrative cost 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 *et al.* [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 *et al.* [2001]). Applying the economics 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 CEO 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 (Department of Justice [2011]), 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 (Department of Commerce[2010], Federal Trade Commission[2010]). However, whether or not regulators and legislators should intervene in the market for privacy is heavily debated in legal (Solove [2004]) and economic (Lenard and Rubin [2005]) circles. Some writers have proposed that U.S. policymakers should rely on self-regulatory frameworks (Mulligan and Goldman [1997]), with targeted regulatory interventions when specific problems – such as data breaches – arise. Such frameworks, however, are often predicated around the idea that consumers can form precise 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)

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

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 potential recipient of, data). According to this perspective, individuals can be relied upon to rationally seek enough privacy to conceal, and to share, the optimal amount of information about themselves.

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 internally consistent as the standard economic perspective assumes. The costs of violations of privacy are often amorphous (e.g., 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 violations are quantifiable because they lead to some tangible cost, the magnitude, timing, and risk of incurring this cost is 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 *et al.* [1990]).

We apply theories from behavioral economics and decision research to investigate the premise that privacy valuations can be precisely estimated. In a series of field and online experiments, we attempt to develop a more nuanced and detailed understanding of privacy valuations, focusing on *informational* privacy and the value individuals assign to the protection of their personal information. The studies we report herein highlight that (policy) discussions about privacy valuations often conflate two different types of transactions 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. To properly understand

privacy decision making, we argue and demonstrate empirically, lessons from behavioral economic research – in particular those arising from the literature on the endowment effect (Kahneman and Tversky [1979]) – must be applied: Personal information is something individuals may feel “endowed” with, in which case one would expect substantially different valuations of the privacy of one’s data depending on whether the focus is on the cost to protect one’s personal information, or the price at which to reveal said information.

Robustly across different scenarios, 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 non-normative influences – that is, factors that, in principle, should not affect decision making. These results challenge the robustness of estimates of privacy valuations proposed in the literature, and also 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. In one of our experiments, for example, 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 didn’t enjoy such belief. In fact, 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] one’s data, in turn, may make individuals feel they have less [more] privacy – and so on.

In addition, we focus on an issue that prior empirical research could have, but has largely failed to, discuss: the *distribution* of privacy concerns across persons. Prior studies that have measured privacy valuations have almost uniformly measured mean valuations (or, at best, minimum and maximum values), without shedding light on the underlying distribution of those valuations. In this paper, we show that privacy valuations are not normally or uniformly distributed, but bipolar, clustering around extreme, focal values. While this might seem to suggest that valuations are relatively stable, this is not the case, because the same individual can easily flip from one extreme of the distribution to the other depending on

contextual factors, including how the values are elicited. Finally, in a series of additional experiments, we test robustness and boundary conditions of our main effects, and examine whether they apply to different types of privacy concerns.

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 [1996], Noam [1996], and 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 *et al.* [2007], and Png *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], and 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] and 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 exact individuals' monetary valuations of privacy. Most of these efforts have focused, either explicitly or implicitly (via the authors' unstated assumptions), on individuals' willingness to accept payment in exchange for disclosing otherwise private information. Huberman *et al.* (2005) used a second-price auction to study the amount of money individuals would require to reveal their weight or height to others. Wathieu and Friedman (2007) showed 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) found significant differences in the price EU citizens would accept to reveal their mobile phone location data, depending on their country of residence. Hui *et al.* (2007) used a field experiment in Singapore to study the value of various privacy assurance measures, finding that privacy statements and monetary incentives could induce individuals to disclose personal information. Chellappa and Sin (2005) also found evidence of a tradeoff between consumer valuation for personalization and concerns for privacy.

Often, this literature has shown that privacy valuations are rather low. For example, Tedeschi (2002) reported on a 2002 Jupiter Research study in which 82% of online shoppers were willing to give personal data to new shopping sites in exchange for the mere *chance* to win \$100. Spiekermann *et al.* (2001) studied subjects' willingness to answer personal questions in order to receive purchase recommendations and discounts, and found that even privacy concerned individuals revealed personal information 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) found that although most survey respondents reported to be concerned about their privacy, only 47% of them would be willing to pay any amount to ensure the privacy of their information. Tsai *et al.* (2011) found that once privacy-relevant information was made salient, participants in an experiment paid moderate premia to purchase both privacy sensitive and non sensitive goods from online merchants with better privacy protection. Varian *et*

al. (2005) and Png (2007) tried to estimate the implicit price that US consumers would pay for the protection from telemarketers, and found values ranging from a few cents to slightly more than \$30.

In the language of economics, the first set of studies focuses on individuals' WTA: the *lowest* price a person would be willing to accept to part with a good (protection of her personal data) she initially owned. The second set of studies focuses on individuals' WTP: the *maximum* price a person would be willing to pay to acquire a good she did not own (protection for her data). 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], Kahneman, Knetsch, and Thaler [1991]) for a vast array of both tangible and intangible goods (see, for instance, Dubourg, Jones-Lee, and Loomes [1994]). Various explanations have been proposed for such WTP/WTA discrepancy (Hanemann [1991]; Hoehn and Randall [1987]), despite valiant attempts at eliminating it (Plott and Zeiler [2005]). By far, loss aversion – the disproportionate weight that people tend to place on losses relative to gains – is the best supported explanation for this WTP-WTA gap (Kahneman and Tversky [1979], Thaler [1980]).

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 above mentioned empirical studies of privacy valuations has explicitly contrasted individuals' willingness to pay to protect data to 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) used conjoint analysis to quantify the value individuals ascribe to website privacy protection, and concluded that “among U.S. subjects, *protection against* errors, improper access, and secondary use of personal information is worth US\$30.49-44.62” (emphasis added). Hann *et al.*'s study is a seminal

contribution in this area: it offered a first insight, and quantification, of the value individuals assign to online privacy, and in doing so it also stimulated 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. Therefore, they cannot determine definitely the “value of protection against errors,” nor the “true” estimate of the value that individuals assign to data - if it was established that those values do differ. A similar problem exists with “revealed preferences” approaches used in other studies: in our experiments, one out of two individuals primed to believe that their privacy was, by default, protected, rejected cash offers for their data; not so, when they had not been thusly primed. Since behaviors can change so radically under the influence of non-normative factors, the preferences they are supposed to reveal also must be mutable and, perhaps, inconsistent.

The distinction between WTP and WTA seems critical in the privacy realm, because real-life, every-day privacy decisions come in both varieties. Analogous to WTP, every day we are faced with opportunities to pay to prevent our personal data from being disclosed – for example, using an anonymous web browsing application, such as Tor,² hides one’s online behavior, but incurs the cost of slower downloads; similarly, deleting cookies to shield, to some extent, one’s browsing habits comes at the cost of having to insert more often registration information across a variety of websites. Analogous to WTA, in other situations we are asked to reveal personal information that we otherwise keep to ourselves, in exchange for some financial benefit – for example, the Internet data company comScore offers its panelists a bundle of products in exchange for monitoring the panelists’ Internet behavior,³ 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 where the items are bought and sold by consumers, and, therefore, objective

² See <https://www.torproject.org/>.

³ See <http://www.comscore.com/>.

prices are formed. In the case of privacy, however, consumers by and large do not participate in (and frequently remain unaware of) 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, making it hard to separate consumers’ valuations of privacy from their valuation of discounts and purchased goods. Trade-offs such as these are, in fact, becoming more common, even inescapable: In some cases, consumers can get access to certain goods or services (such as listening music on Spotify,⁴ or commenting on news stories on the LA Times website), only through a social network that tracks their behavior and links it to their actual identities (Facebook).

As a result, managers and policy makers in search of guidelines to assess the value citizens and consumers assign to the protection of their data must rely on estimates from research studies, or anecdotal revealed preferences arguments. The need therefore arises for carefully scrutinizing the assumptions those estimates are based upon, and vetting the robustness of the predictions based on those assumptions. Our research shows that privacy valuations are very sensitive to non-normative factors: We may not be willing to spend even a few cents to protect a certain piece of data, and yet we may reject offers of several dollars to sell the same data. Which one of these valuations would be the “true” value of the privacy of our data? Both cannot simultaneously reflect our true preferences.

⁴ See <http://www.spotify.com/>.

The dichotomy between WTP and WTA is just one example of the notion that preference for privacy may not only be context-dependent, but malleable and uncertain, and suggests that ordinary studies investigating privacy valuations may not tell us much about whether consumers will actually pay to protect their data. Behavioral economists have highlighted that non-normative 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 (such as hyperbolic discounting – see Acquisti [2004]; or the illusion of control – see Brandimarte *et al.* [2010]) on privacy decisions, and how those decisions deviate from the abstractly rational path predicated on neoclassical economic theory.

In the following sub-section, 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 therefore the extent to which researchers are able measure those valuations.

2.3 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: $u(w+WTA,p^-) = u(w,p^+)$. 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: $u(w-WTP,p^+) = u(w,p^-)$. The implication of these equations is that WTA will not necessarily be identical to WTP, and specifically, if privacy is a normal good that becomes valued more as one becomes wealthier, it is possible that $WTA > WTP$, 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 WTA/WTP discrepancy 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 tradeoffs 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, leading 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, leading 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 non-normative 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 PGP encryption (which protects email 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 hypothesize that:

(Hypothesis 1) Willingness to pay and willingness to accept 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, $u(w^+, p^-) > u(w^-, p^+)$, simply depending on how the question is framed. This would suggest that: 1) the minimum price a consumer will be willing to accept to allow her data to be revealed may be higher than the maximum price she will be willing to pay to avoid having her 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; and more broadly, 2) privacy preferences, while not arbitrary, are malleable to non-normative factors, and can be internally inconsistent, in that the same cash-for-data offer may be accepted or rejected for non-normative 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 based on contextual cues with little normative justification. Consider, specifically, 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 *et al.* [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 privacy protective 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.*

Finally, empirical research on privacy, to date, has examined mean valuations across individuals, but not the *distribution* of those valuations. Uncertainty and ambiguity associated with privacy trade-offs, coupled with the idiosyncrasy of privacy concerns, may again produce unusual distributions of privacy valuations. For instance, John et al. (2011) have found that individuals, in the absence of environmental cues that trigger privacy concern, fail to take privacy risks into account in their decision making. This leads to some surprising effects – for example, assurances of anonymity can, contrary to their typical purpose, cause people to ‘clam up’ and resist sharing information because they trigger, but do not fully allay, concern about privacy. The idea that people either under- or over-react to threats to privacy would also be consistent with the probability weighting function of Prospect Theory (Kahneman and Tversky [1979]) according to which people either tend to ignore, or overweight, outcomes associated with small probabilities. Furthermore, if people don't ordinarily think about privacy, but when they do tend if anything to overweight it, then, contrary to the usual finding that valuations of goods tend to be normally distributed across people, valuations of privacy could, instead, be better approximated by a bimodal distribution. Based on this line of thinking, we conjectured that, unlike ordinary private goods,

Non-normality of valuations: *Privacy valuations are not normally distributed.*

3. THE EXPERIMENTS

We tested our hypotheses in a series of experiments that shared a common design. Subjects were asked to choose between gift cards that varied with respect to their privacy features and monetary values. Across all experiments, we operationalized informational privacy concerns as 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 a) their initial endowment and b) the order in which choices were presented.

Experiment 1 tested Hypotheses 1 and 2 in the field, with real gift cards. Experiment 2 was a hypothetical survey that replicated the results of Experiment 1, but enabled us to examine individual privacy valuations and test whether, as conjectured, they are bimodally distributed. Experiments 3a-d were hypothetical choice follow-up studies that tested robustness and boundary conditions for the findings of the prior two Experiments.

3.1 Experiment 1: Endowment and order effects

Experiment 1 was a field experiment in which subjects were offered real VISA gift cards that could be used to purchase goods from any online or offline store where debit cards are accepted. Shoppers at a shopping mall 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 distract attention from) the gift card that subjects were given as 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 either endowed with the \$10 anonymous card or the \$12 identified card, before being offered the option to swap one card for the other. Those 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 appeared first, and in the other condition the identified \$12 card appeared first). 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, thus forming 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 below:

1. [*\$10 Endowed*] *Keep the anonymous \$10 card or exchange for an identified \$12 card*
2. [*\$12 Endowed*] *Keep the identified \$12 card or exchange for an anonymous \$10 card*
3. [*\$10 Choice*] *Choose between an anonymous \$10 card (appearing first) and an identified \$12 card*
4. [*\$12 Choice*] *Choose between an identified \$12 card (appearing first) and an anonymous \$10 card*

Note that all subjects in the first four conditions, regardless of the condition to which they had been randomly assigned, faced the exact same alternatives: a \$10 anonymous card or a \$12 identified card. However, the gift card endowment in two of the conditions generated a different framing of the choice faced by the subjects: for those in the [*\$10 Endowed*] conditions, 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*] conditions, the question was framed as an implicit choice to pay \$2 in order to avoid having one’s future purchase data made available to the researchers.² 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/WTA dichotomy, and/or order effects.³

3.1.1 Procedure

The experimental procedure is summarized in this section (complete details are available in the online Appendix). Experiment 1 took place on three weekend days at a Pittsburgh shopping mall. Female

research assistants stood at the entrance of two women's clothing stores and approached female shoppers as they entered, asking them to complete a brief survey. To make the decoy survey realistic, shoppers were told that the survey was designed to assess people's attitudes toward spending money. Interested shoppers were given a coupon valid for a gift card upon completion of a short survey. After completing the survey and upon exiting the store, each subject gave her coupon to the experimenter, who then asked the subject (regardless of the 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.

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. In the *choice* conditions, subjects were only presented with one description sheet that listed and described both cards, one after the other, with order of description presentation manipulated between-subjects. Subjects then indicated which card they would like and selected their card from the appropriate bin. Subjects were then asked to provide their email address.

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

3.1.2 Results

Three-hundred and forty-nine female subjects participated in the study (M age=35; Median age=35; 83.6% Caucasian; all not significant between conditions). Upon exiting the store, the majority (92.3%) of subjects returned to the experimenter to redeem their gift card coupon. Subjects were more likely to redeem their coupon if they completed the survey upon entry (95.4%) versus upon exiting the store (88.9%) (χ^2 (1) = 5.14, p = 0.023). However, the likelihood of completing the survey upon entry versus exit did not differ between conditions (χ^2 (4) = 3.71, p = 0.447), nor did redemption rates (χ^2 (4) = 2.35, p = 0.673).

Gift card choice. Virtually everyone in the “rationality check” control condition (95.7%) selected the \$12 anonymous card, suggesting that subjects understood and took the task seriously. This condition is excluded from the rest of the analyses.

The proportion of people choosing the \$10 anonymous card was highest when subjects had been endowed with it (52.1%); followed by the choice condition in which the \$10 card was listed first (42.2%); followed by the choice condition in which the \$10 card was listed second (26.7%); and lowest (9.7%) for those endowed with the \$12 identified card (see Figure 1). Subjects in the endowed conditions displayed a tendency to keep the card they had been endowed with – confirming previous results on the power of default settings on privacy decision making (Johnson *et al.* [2002]). However, and more interestingly, while 90.3% of subjects in the \$12 endowed condition kept the \$12 card, only 52.1% 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 (χ^2 (1) = 27.24, p < 0.001). More importantly, a majority of subjects in the \$10 endowed condition (52.1%) rejected an offer of \$2 (WTA) to switch to an identified card in exchange for giving away their future purchase data. However, only a small minority of subjects (9.7%) paid two dollars for privacy (WTP), by switching from the \$12 identified card to the \$10 anonymous card to protect the same data.

The two choice conditions – differing only in the order in which the cards were described– are marginally significantly different from each other ($\chi^2(1) = 3.64, p = 0.056$): subjects seemed more likely to choose the card that was described first. Specifically, when the \$12 identified card was listed first, 73.3% of subjects chose it, whereas when it was listed after the description of the \$10 anonymous card, only 57.8% 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 (1) over the privacy enhanced \$10 gift card (0).⁴ We ran one regression for the two endowed conditions (second column) and one for the two choice conditions (third column). We used a dummy variable (*\$10Card*) to control for which card the subject was endowed with (or presented first): the \$10 card (1) or the \$12 card (0). Both models are significant. In the endowed conditions, *\$10Card* is strongly significant and negative ($p < 0.001$): subjects endowed with a \$10 card were less likely to choose to give away their data for \$2. This result strongly supports Hypothesis 1. In the choice conditions, *\$10Card* is negative and weakly significant ($p = 0.1$), providing mild support for Hypothesis 2 (presenting a privacy enhanced option before the less privacy enhancing one sends a signal that the former is inherently more valuable), but also indicating that order effects are less strong than endowment effects.

[Figure 1 about here]

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%) had used their cards. Subjects who had chosen the more valuable card were slightly more likely to have used it (90.7% of those with \$12 cards versus 81.8% of those with \$10 cards; Pearson $\chi^2(1) = 4.25, p = 0.039$). There were no significant differences in the propensity to use the card depending on the initial conditions of assignment (whether the subject had been *initially* endowed with, or had to initially

choose, a card; Pearson $\chi^2(1) = 0.16$, $p = 0.688$), or whether the subject had been initially assigned an anonymous or identified card (Pearson $\chi^2(1) = 1.28$, $p = 0.258$).

We investigated whether subjects used their cards at different types of stores, depending on card identifiability. Stores were classified as potentially privacy sensitive (e.g. lingerie stores such as “Victoria’s Secret”) or not (cafes, convenience stores, supermarkets). We found modest anecdotal evidence of differences in store patronage depending on card identifiable. For instance, all of the eight purchases recorded at Victoria’s Secret were completed with the more valuable but less privacy protected card. This evidence should be considered as merely suggestive: store selection was not designed as part of the controlled experiment, since subjects could use their cards at any online or offline store.

[Table 1 about here]

Subject’s decision making. In the exit questionnaire, we asked subjects to explain why they choose one card over the other. Explanations provided by subjects who chose the \$10 card often referenced privacy concerns, and specifically 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 ... 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 the following concepts: “More money to spend! ... Because it was more money!” or even referred specifically 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 ... I don’t mind if you know where I spend my money.”

Analysis. In Experiment 1, subjects across experimental conditions chose gift cards in different proportions merely depending on the framing of the choice. In doing so, they implicitly assigned, and revealed, dramatically different values to the privacy of their data. Valuations in the two endowed conditions were different from the choice conditions, and the valuations in the choice conditions differed

based on which option was presented first. For example, more than half of subjects in the anonymous \$10 endowed condition rejected an offer of \$2 to reveal their future purchase data (that is, an increase of 20% 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. Within the context of the experiment, their WTA was therefore larger than (or at best equal to) \$2. Evidently, these subjects felt endowed with the protection of their information (naturally, this is not an absolute statement about the subjects' universal privacy preferences: the \$2 amount is itself function of various factors held constant across the experiment's conditions, including – for instance – switching costs). By contrast, fewer than 10% of subjects endowed with the identified \$12 card chose to give up \$2 (a 17% decrease in their initial endowment) to protect future purchase data. The overwhelming majority of those subjects refused to pay \$2 to protect their future purchase data – they decided that \$2 was *too much* to protect their privacy. These results imply that subjects were *five times* more likely to choose privacy in one condition over the other, even though all subjects faced exactly the same choice. These patterns stand in contrast to results in the literature purporting to measure objective, true valuations of privacy.

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 ($0.52 * \$2 + 0.48 * \0), and that in the [\$12 Endowed] condition at roughly 19 cents. This represents a WTA/WTP ratio of 5.47 – markedly larger than the average ratio observable 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 private 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,

making the estimation of privacy valuations for the benefits of public policy and decision making even more challenging.

The results challenge the reliance on “revealed preferences” arguments to conclude, based on the fact that few users take advantage of available protective solutions, that they do not care for privacy (Rubin and Lenard [2002]). In our experiment, 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 was, by default, protected ([$\$10$ Endowed] condition), than when they believed that their data would be, by default, revealed ([$\$12$ Endowed] condition). The latter condition is arguably more likely to reflect consumers’ actual beliefs and fears about the current state of privacy protection (surveys repeatedly find that most U.S. residents do not think their privacy is adequately protected; see, for instance, Kelsey and McCauley [2008]). Experiment 1 therefore suggests that when consumers feel that their privacy is protected, they value it much more than when they feel their data has already been, or may be, revealed. In fact, even just making subjects *feel* that they start from a position in which their data, by default, is not protect, is sufficient to decrease the value they assign to the privacy of that data – thus making the expectation (one’s data will not be protected) more likely to become reality (one’s data is traded away).

3.2 Experiment 2: The distribution of privacy valuations

Experiment 2 was a two-part survey-based experiment. In the first part, subjects were asked to imagine receiving a gift card as payment for participating in a research study. After reading about the value and the characteristics of the card, subjects were asked whether they would like to exchange it for a card of different value and with different privacy features. This first part was similar to Experiment 1, but differed in that subjects – depending on the experimental condition - were asked to choose between $\$10$ cards with privacy, and $\$12$ or $\$14$ cards without privacy. Hence, Experiment 2 allowed us to test whether the WTP/WTA dichotomy found in Experiment 1 would extend to cases where the differential cash value in the card was larger than $\$2$. Beyond this difference, Experiment 2 also included a second part, the purpose of which was to estimate subjects’ *distributions* of privacy valuations. After stating their card

choice, subjects were presented with follow-up choices, based on increasing or decreasing differences in the values of the card, and were asked to repeat their selections.

3.2.1 Procedure

The experiment was a 2x2 between-subjects factorial design. Subjects were randomly assigned to experimental conditions that differed by the type of card they were initially offered. We manipulated a) whether subjects were (hypothetically) initially endowed with a trackable (WTP) or an untrackable card (WTA), and b) the difference in the value between the two cards (trackable card worth \$2 or \$4 more than untrackable card). We refer to conditions in which subjects were assigned a trackable card as “WTP” since they relate to the question of how much (if anything) subjects would be willing to pay back to protect their data, and conditions in which subjects were assigned an untrackable card as “WTA” since they relate to the question of how much (if anything) subjects would be willing to accept to give away their data. The tradeoff in each of the four conditions was as follows:

1. *[WTA/Δ2] Keep \$10 card which cannot be tracked, or exchange for \$12 card which will be tracked*
2. *[WTA/Δ4] Keep \$10 card which cannot be tracked, or exchange for \$14 card which will be tracked*
3. *[WTP/Δ2] Keep \$12 card which will be tracked, or exchange for \$10 card which cannot be tracked*
4. *[WTP/Δ4] Keep \$14 card which will be tracked, or exchange for \$10 card which cannot be tracked*

In addition, we used a fifth condition ([WTA/Δ2 Control]) to test whether subjects may be sensitive to slight changes in the description of the cards. In this condition, subjects were asked to choose between keeping the \$10 card which cannot be tracked (as in condition [WTA/Δ2]), or exchange it “for the \$12 card which *may* be tracked” (emphasis added).

The first page of the questionnaire stated that there were two types of gift cards: trackable and untrackable (Appendix C). Purchases made with a trackable card would be “tracked by researchers” and “linked to the name of the participant.” Purchases made with an untrackable card would “not be tracked by researchers” and therefore would “not be linked to the name of the participant.” Subjects were then asked whether they would like to keep the card they were initially offered, or exchange it for the other card. After answering this question, subjects were instructed to turn the page and answer the follow-up questions that allowed us to estimate their distribution of privacy valuations. On the final page of the questionnaire, subjects answered demographic questions.

3.2.2 Results

Experiment 2 was run at cafeterias in hospitals in Pittsburgh. Subjects were recruited on site; each was offered a chocolate bar for completing the questionnaire. Two hundred and forty subjects participated in the study (46.2% female; M age=33; Median age=35; 75.0% Caucasian); each was randomly assigned to one of the five experimental conditions (50 subjects participated in condition [WTA/ Δ 2], 45 in condition [WTA/ Δ 4], 51 in condition [WTP/ Δ 2], 44 in condition [WTP/ Δ 4], and 50 in the [WTA/ Δ 2 Control] condition). Except for a slight overrepresentation of females in Condition [WTA/ Δ 2], there were no other significant demographic differences between conditions (we did not find any gender effect on card choice).

In the conditions in which we asked subjects to choose between a \$10 anonymous card and \$12 trackable card (conditions [WTA/ Δ 2] and [WTP/ Δ 2]), we found, as hypothesized, a significant effect of card endowment on card choice.⁵ When endowed with the \$10 untrackable card, 60.0% of subjects claimed they would keep it; however, when endowed with the \$12 trackable card, only 33.3% of subjects claimed they would switch to the untrackable card ($\chi^2(1) = 6.76, p = 0.009$). We found a similar pattern in the conditions in which we asked subjects to choose between a \$10 anonymous card and a \$14 trackable card (conditions [WTA/ Δ 4] and [WTP/ Δ 4]): 60.0% of subjects endowed with the \$10 card claimed they would keep that card, but only 41.5% of the subjects endowed with the \$14 card indicated

that they would switch to the \$10 card. In this case, however, the difference was only marginally significant ($\chi^2(1) = 2.95, p = 0.086$).

[Table 2 about here]

To control for age and gender effects, we ran logistic regressions on the binary choice variable using a probit model. We included data from the four comparable conditions and regressed age, gender, and dummy variables representing the conditions over a dichotomous dependent variable, representing the selection of the traditional gift card (1) over the privacy enhanced gift card (0) (see Table 2). We used one dummy variable to control for the conditions which contrast \$10 and \$12 cards ($\Delta 2=1$) versus \$10 and \$14 cards ($\Delta 2=0$), and another dummy to control for the conditions in which the subjects were endowed with the untrackable card and were offered to accept more money to switch to the tracked card ($WTA=1$). Age is a discrete variable and gender is a binary dummy (1=female). The model is significant, and the WTA/WTP effect is strongly significant: subjects in the WTA conditions are much less likely to switch to the trackable cards than subjects in other conditions. These results are consistent with those of Experiment 1, and show that the endowment effect extends to larger value differences across the card than those examined in Experiment 1.

However, and importantly, we found no effect of the difference in card values (i.e. $\Delta \$2$ vs. $\Delta \$4$) on subjects' card choice. We also found that the interaction between card value and endowment is not significant (last column in Table 2). In fact, there was no difference in the percentage of subjects who kept the untrackable \$10 card when offered to exchange it for a \$12 or a \$14 trackable card (in both cases, 60.0% of subjects claimed they would keep it; Pearson $\chi^2(1) = 0.00, p = 1$). Similarly, there was no significant difference in the number of people who claimed they would switch to a \$10 untrackable card *from* a \$12 or \$14 trackable card (33.3% in the former case, and 41.5% in the latter case claimed they would switch; Pearson $\chi^2(1) = 0.91, p = 0.339$). These results suggest that privacy valuations, within the context of the experiment, did not vary significantly in the [\$2-\$4] interval. For instance, some

individuals may have valued privacy protection a lot (\$4 or more, so their choice would not change depending on whether they are offered \$2 or \$4 for their data); other individuals may have valued such protection barely at all (less than \$2, so being offered \$2 or \$4 would not make a difference to them either); but very few individuals valued the privacy of the purchase data between \$2 and \$4. This result suggests that privacy valuations are not uniformly or even normally distributed, but instead, clustered around focal points. The follow-up questions in the second part of Experiment 2, which were designed to elicit a distribution of privacy valuations, allowed us to examine the issue directly.

The distribution of privacy valuations. The follow-up questions in Experiment 2 focused on whether the subject would have chosen the same or an alternative card if the values of those cards had been different. The alternative values presented in the follow-up questions depended on the subject's card choice as specified on the first page, and incremented (or decremented) by as little as 25 cents or as much as a few dollars (see Appendix C). Based on the responses to the follow-up questions, we constructed a variable representing "brackets" of privacy valuations – the approximate monetary range that individuals assigned to the untrackable card. For instance, consider the subjects who chose to keep a \$10 untrackable card (rather than switching to a \$12 trackable card). We define their "privacy valuation" to be at least \$2 (once again, we note that this is not an absolute statement about the subjects' universal privacy preferences, as the various amounts are themselves function of other factors held constant across the conditions, such as switching costs). Suppose that the same person then indicated that she would have also kept the untrackable card if it had been worth \$9, but *not* if it had been worth \$8. We would then infer a (self-reported) valuation for the privacy of her purchase data to be *at least* \$3 (the difference between the offered \$12 and the hypothetically endowed \$9), but *less than* \$4 (the difference between the offered \$12 and the hypothetically endowed \$8). We then took the lower boundary of each bracket, and constructed the histograms presented in Figure 2 (for instance, if the subject's valuation was calculated to lie within the 0c to 0.25c bracket, we used a value of 0 for the histogram; if it was between 0.50 and 0.75, we used 0.50; and so forth).

Figure 2 presents brackets of values for each of the five experimental conditions, as well as the values aggregated across conditions (bottom right quadrant). Consistent with Conjecture 1, all distributions (with the exception of the WTP/ $\Delta 2$ condition) are bimodal (also, consistent with Hypothesis 1, the bimodality is more accentuated in the conditions in which subjects were endowed with the privacy enhanced card). We used non-parametric rank sum Mann-Whitney tests to compare the distributions of valuations across conditions, and found statistically significant differences when contrasting the two \$10 vs. \$12 conditions ($z = 3.67$, $p < 0.001$) and the two \$10 vs. \$14 conditions ($z = 2.647$, $p = 0.008$).

In both cases, the conditions endowed with the more valuable but unprotected card tend to assign less value to the privacy enhanced card, which is consistent with Hypothesis 1 and the results presented in Section 3.1. The modal valuation is one of the extreme points for all five conditions (specifically, it is “between 0 and 25 cents” for three conditions, and “larger than \$11” for the other two); the *second* modal valuation is the *other* extreme for four out of five conditions.⁶ Shapiro-Wilk, Shapiro-Francia, and Skewness-Kurtosis tests on the bracket data all strongly rejected the hypothesis of normality of distribution of valuations ($p < 0.05$ within each condition). Hartigan and Hartigan (1985)’s dip test for unimodality also rejected the hypothesis of unimodality (and uniformity) for conditions [WTA/ $\Delta 2$] and [WTA/ $\Delta 4$] and the Control condition ($p < 0.001$), implying bimodality, and was borderline for the [WTP/ $\Delta 4$] condition ($p = 0.11$). It was not rejected, however, for condition [WTP/ $\Delta 2$] ($p = 0.26$), where the lowest possible valuation was the dominant choice for most subjects.

[Figure 2 about here]

Falsification tests. As a falsification test of the bimodality result, we ran a new battery of surveys using the exact same language in the [WTA/ $\Delta 2$] and [WTP/ $\Delta 2$] conditions. In this new set of surveys, we first asked subjects to hypothetically choose between a \$10 gift card *plus a physical good*, and a \$12 card *with no such good*. In other words, we applied our experimental design to a scenario where WTP and WTA were estimated for an ordinary private good, instead of privacy. Next, following the design of Experiment 2, we posed follow-up questions to estimate the distribution of valuations of the goods. Specifically, in

separate falsification tests, we measured subjects' valuations for three goods whose average eBay price fell in the \$2 to \$3 range: an eraser, a desktop aluminum calendar, and an IKEA umbrella. At least 80 subjects were recruited online and used for each falsification test.

When testing WTP and WTA for these physical goods using the design of Experiment 2, the bimodality of the distributions disappears. For example, consider Figure 3: the left quadrant represents the aggregate distribution of *privacy* valuations, combining the familiar results of Experiment 2's conditions [WTA/ $\Delta 2$] and [WTP/ $\Delta 2$]; the bimodality is readily apparent. The right quadrant represents the aggregate distribution of valuations for an IKEA umbrella, as determined from the subjects' choices between a \$10 card and an IKEA umbrella or a \$12 card without such umbrella (n=82). The distribution is no longer U-shaped, but skewed and unimodal ($\text{diptest}_{[\text{WTP/umbrella}]}: p = 0.28$; $\text{diptest}_{[\text{WTA/umbrella}]}: p = 0.10$).

[Figure 3 about here]

3.3 Follow-up WTP/WTA Experiments

We conducted four additional experiments to test the robustness and boundary conditions of the WTP/WTA privacy gap observed in Experiments 1 and 2. Experiment 3a tested whether the effects in Experiments 1 and 2 were unique to privacy features included on a gift card, or whether other gift card features (such as convenience) also elicit endowment effects. Whereas Experiment 1 and 2 pertained to informational privacy (i.e. concerns over the treatment of one's personal data, as operationalized by purchase information; see Tsai *et al.* [2011]), in Experiment 3b we tested whether the endowment effect would extend to another type of privacy concern (namely, location privacy; see Cvrcek *et al.* [2006]). Finally, Experiments 3c and 3d tested boundary conditions of the endowment effect in privacy valuations.

All subjects in the follow-up experiments were recruited from an online platform (MTurk) managed by Amazon.com. Subjects were offered a small fixed payment to participate in a "short online survey." They had to be at least 18 years old and have an "approval rate" (based on their history of tasks completed on the platform) of at least 99% to participate in the study. Subjects who chose to take the

survey were randomly assigned to one of eight experimental conditions (two conditions each for Experiment 3a, 3b, 3c, and 3d). Summary results of these experiments are presented in the rest of this section. For brevity, we only report the results of χ^2 tests of the proportion of gift cards chosen across the conditions.

3.3.1 Experiment 3a

Experiment 3a tested whether the results of Experiment 2 are unique to the privacy features of a gift card. Instead of choosing between differently-valued gift cards with various privacy features, subjects had to choose between cards of different value and *convenience*. Subjects in Condition 1 ($n = 45$) were asked to imagine they had received a \$10 card that could be used both in stores inside the mall and online. Then, they were offered a more valuable but less convenient \$12 card that could only be used in stores inside the mall. Subjects in Condition 2 ($n = 57$) faced the reverse proposition (we kept the values of the card at \$10 and \$12 in order to make the results directly comparable to Experiment 2). Unlike in the case of privacy (that is, Experiment 2), no endowment effect was found: both in Condition 1 (57.8%) and in Condition 2 (66.7%), the relative majority of subjects opted for the more valuable card ($\chi^2(1) = 0.85, p > 0.3$), even though a significant portion of subjects in both conditions found the more convenient card desirable (more than 40% in Condition 1 and more than 30% in Condition 2). Experiment 3a therefore suggests that not all intangible gift card features elicit the same endowment effect that we observed for privacy.

3.3.2 Experiment 3b

Experiment 3b was designed to establish whether the endowment effect found in our experiments would hold for a different form of privacy concern – namely, location privacy (Cvrcek *et al.* [2006]). Similar to Experiment 2, subjects were asked to imagine having received a small payment and a VISA gift card for participating in a research study. However, instead of choosing between different valued gift cards that would or would not allow researchers to track *purchases* made with that card (as in Experiment

2), subjects in Experiment 3b had to choose between: 1) a \$18 card, on the condition that their *location* would be recorded by the researchers for one day via the subjects' cell phone GPS; and 2) a \$12 card, which came with no such condition. Therefore, the trackable card was \$6 more valuable than its alternative. We chose a larger monetary differential than in Experiments 1 and 2 because we predicted that subjects would consider location data to be more sensitive than purchase data.

In Condition 1, subjects (n=78) were asked to imagine they had been given the \$12 card. Then, they were offered the option of changing to the \$18 card. In Condition 2, subjects (n=77) were asked to imagine they had been given the \$18 card. Then, they were the \$12 card. The results confirmed the existence of a privacy endowment effect: in Condition 1, 64.1% of subjects chose to keep the \$12 card, but only 35.9% of subjects made that decision in Condition 2 ($\chi^2(1) = 13.07, p < 0.001$). Experiment 3b therefore suggests that the endowment effect also arises for the valuation of physical privacy, in addition to informational privacy valuations examined in Experiments 1 and 2.

3.3.3 Experiments 3c and 3d

The design of Experiments 3c and 3d was identical to Experiment 2: subjects had to choose between two differently-valued gift cards that would or would not allow the researchers to track purchases made with the card. However, to test the boundary conditions of the privacy endowment effect, we varied the differential in the values of the two cards, making it very large - \$15 - in Experiment 3c, and very small -- 25 cents in Experiment 3d.

Subjects in Experiment 3c, Condition 1 (n = 52) were asked to imagine they had been given a \$10 privacy-enhanced card. Then, they were offered the option of replacing it with a \$25 card without privacy protection. Subjects in Condition 2 (n = 41) faced the reverse choice (from \$25 to \$10). In both conditions, a majority of subjects chose the more valuable (but less private) card (69.2% in Condition 1 and 80.5% in Condition 2). The endowment effect was no longer significant ($\chi^2(1) = 1.52, p > 0.2$): when the difference in the two gift cards' values is too large, most subjects choose the most valuable card

because the monetary advantage trumps the privacy concerns. This confirms the existence of an upper ceiling to individuals' privacy valuations.

Subjects in Experiment 3d, Condition 1 ($n = 51$), were asked to imagine receiving a \$10 privacy-enhanced card. They were then offered the option of replacing it with a \$10.25 card without privacy protection. A majority (72.5%) kept the privacy enhanced card. Subjects in Condition 2 ($n = 46$) faced the reverse choice (from \$10.25 to \$10); of them, 52.2% kept the \$10.10 card, and 47.8% chose the privacy enhanced card. While the endowment remains significant at the 5% level ($\chi^2(1) = 6.10, p = 0.013$), the results indicate that when the monetary benefit of giving away one's purchase data is so small (25 cents), the \$10 card becomes an appealing option also for close to half of subjects in the "WTP" condition (that is, those endowed with the \$10.25 card). In fact, the proportion of subjects that replaced a more valuable card with a less valuable but privacy enhanced card was significantly larger in Experiment 3d (47.8% when delta between the two cards is 10 cents) than in Experiment 3c (19.5% when the delta between the two cards is \$15) ($\chi^2(1) = 7.69, p = 0.006$).

4. CONCLUSIONS

In everyday life, individuals face privacy trade-offs of two types, which 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. The experiments just presented empirically shown the impact of this difference on privacy valuations; those who could give up privacy assurances for a material benefit made decisions that suggested much greater concern for privacy than those who could, in effect, purchase greater privacy at a price. We also observed strong order effects, reinforcing the conclusion that privacy preferences are not precisely defined. Finally, we find that privacy valuations are not normally or uniformly distributed, but bimodal, clustered around extreme, focal values.

These findings have implications for both empirical and theoretical economic analyses of privacy. At an empirical level, our findings should caution against the uncritical use of valuations of privacy that have used single methods – e.g., *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 her 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 valuation 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 importantly, this research raises the issue of individuals’ abilities to rationally navigate issues of privacy. From choosing whether or not to join a grocery loyalty program, to posting embarrassing personal information on a public website, individuals constantly make privacy-relevant decisions which impact their well-being, and this research suggests that they do so inconsistently.

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 towards 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 (e.g., Rubin and Lenard [2002]), on the grounds that 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, A. 2004. "Privacy in Electronic Commerce and the Economics of Immediate Gratification." *Proceedings of ACM Electronic Commerce Conference (EC '04)*. New York, NY: ACM Press, 21-29.
- Acquisti, A. and H. Varian, 2005. "Conditioning Prices on Purchase History." *Marketing Science*, 24(3), 1-15.
- Barnes, R., 2012. "Supreme Court: Warrants Needed in GPS Tracking." *Washington Post*, January 23.
- Brandimarte, L., A. Acquisti, and G. Loewenstein, 2010. "Privacy Concerns and Information Disclosure: An Illusion of Control Hypothesis." *Conference in Information Systems and Technology (CIST)*.
- Brookshire, D.S., R.C. d'Arge, W.D. Schulze, and M.A. Thayer, 1981. "Experiments in valuing public goods." In: V.K. Smith (ed) *Advances in Applied Microeconomics: Volume 1*, Greenwich CT: JAI Press.
- Calzolari, G. and A. Pavan, 2006. "On the Optimality of Privacy in Sequential Contracting." *Journal of Economic Theory*, 130(1), 168-204.
- Chaum, D. 1983. "Blind signatures for untraceable payments." *Advances in Cryptology - Crypto '82*, Springer-Verlag, 199-203.
- Chellapa, R. and R.G. Sin, 2005. "Personalization Versus Privacy: An Empirical Examination of the Online Consumers' Dilemma." *Information Technology and Management*, 6(2-3), 181-202.
- Culnan, M. J. and P.K. Armstrong, 1999. "Information Privacy Concerns, Procedural Fairness, and Impersonal Trust: An Empirical Investigation." *Organization Science*, 10(1), 104-115.
- Cvrcek, D., M. Kumpost, V. Matyas, and G. Danezis, 2006. "A Study on the Value of Location Privacy." *Proceedings of Workshop on Privacy in the Electronic Society (WPES '06)*, 109-118.
- Department of Commerce, 2010. "Commercial Data Privacy and Innovation in the Internet Economy: A Dynamic Policy Framework." Internet Policy Task Force Green Paper.
- Department of Justice, 2011. "Identity Theft Reported by Households, 2005-2010." Bureau of Justice Statistics, Crime Data Brief.
- Dinev, T. and P. Hart, 2006. "An Extended Privacy Calculus Model for e-Commerce Transactions." *Information Systems Research*, 17(1), 61-80.
- Dubourg, W.R., M.W. Jones-Lee, and G. Loomes, 1994. "Imprecise preferences and the WTP-WTA disparity." *Journal of Risk and Uncertainty*, 9(2), 115-133.
- Federal Trade Commission, 2010. "Protecting Consumer Privacy in an Era of Rapid Change." FTC Report.
- Gonsalves, A. 2010. "Facebook CEO: Less Privacy is Social Norm." *InformationWeek*, January 12.

- Hammack, J. and G.M. Brown, 1974. *Waterfowl and Wetlands: Toward Bioeconomic Analysis*. Baltimore, Maryland: John Hopkins University Press.
- Hanemann, M.W., 1991, "Willingness to Pay and Willingness to Accept: How Much Can They Differ?" *American Economic Review*, 81, 635-647.
- Hann, I.H., K. L.Hui, T. Lee, and I. Png, 2007. "Overcoming Information Privacy Concerns: An Information Processing Theory Approach." *Journal of Management Information Systems*, 24(2), 13-42.
- Hartigan, J. A. and P. M. Hartigan, 1985. "The Dip Test of Unimodality." *Annals of Statistics* 13(1), 70-84.
- Hirshlerifer, J., 1980. "Privacy: Its Origins, Function and Future." *Journal of Legal Studies*, 9, 649.
- Hoehn, J.P. and A. Randall, 1987. "A Satisfactory Benefit Cost Indicator from Contingent Valuation." *Journal of Environment, Economics and Management*, 14, 226-247.
- Horowitz, J.K. and K.E. McConnell, 2002. "A Review of WTA / WTP Studies." *Journal of Environmental Economics and Management*, 44, 426-244.
- Huberman, B., E. Adar, and L. Fine, 2005. "Valuating Privacy." *IEEE Security & Privacy*, 3(5), 22-25.
- Hui, K.-L., H-H. Teo, S.-Y. Lee, 2007. "The Value of Privacy Assurance: An Exploratory Field Experiment." *MIS Quarterly*, 31(1), 19-33.
- Kahneman, D., 1986. "Valuing Environmental Goods: An Assessment of the Contingent Valuation Method." In: *Valuing Environmental Goods: An Assessment of the Contingent Valuation Method*, R. Cummings, D. Brookshire, and W. Schulze (eds), Totowa, NJ.
- Kahneman, D., J.L. Knetsch, and R. H. Thaler, 1990. "Experimental Tests of the Endowment Effect and the Coase Theorem." *Journal of Political Economy*, 98(6), 1325-1348.
- Kahneman, D., J.L. Knetsch, and R. H. Thaler, 1991. "Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias." *Journal of Economic Perspectives*, 5(1), 193-206.
- Kahneman, D. and A. Tversky, 1979. "Prospect Theory: An Analysis Of Decision Under Risk." *Econometrica*, 47(2), 263-292.
- Kelsey, J. and M. McCauley, 2008. "Consumer Reports Poll: Americans Extremely Concerned About Internet Privacy - Most Consumers Want More Control Over How Their Online Information Is Collected and Used." Consumerunion.org, September 25.
- Knetsch, J.L., 1989. "The Endowment Effect and Evidence of Nonreversible Indifference Curves." *American Economic Review*, 79(5), 1277-1284.
- Laudon, K.C., 1996. "Markets and Privacy." *Communications of the ACM*, 39(9), 92-104.
- John, L.K., A. Acquisti, and G. Loewenstein, 2011. "Strangers on the Plane: Context Dependent Willingness to Divulge Personal Information." *Journal of Consumer Research*, 37(5), 858-873.

- Johnson, E., S. Bellman, and G. Lohse, 2002. "Defaults, Framing and Privacy: Why Opting In-Opting Out." *Marketing Letters*, 13(1), 5-15.
- Laufer, R.S. and M. Wolfe, 1977. "Privacy as a Concept and a Social Issue: A Multidimensional Developmental Theory." *Journal of Social Issues*, 33(3), 22-42.
- Mulligan, D. and J. Goldman, 1997. "The Limits and the Necessity of Self-Regulation: The Case for Both." In: *Privacy and Self-Regulation in the Information Age*, National Telecommunications and Information Administration, US Department of Commerce.
- Noam, E.M., 1996. "Privacy and Self-regulation: Markets for Electronic Privacy." In: *Privacy and Self-Regulation in the Information Age*, National Telecommunications and Information Administration, US Department of Commerce.
- Plott, C.R. and K. Zeiler, 2005. "The Willingness to Pay/Willingness to Accept Gap, The 'Endowment Effect,' Subject Misconceptions and Experimental Procedures for Eliciting Valuations." *American Economic Review*, 95(3) 530-545.
- Png, I.P.L., 2007. "On the Value of Privacy from Telemarketing: Evidence from the 'Do Not Call' Registry." Working paper, available at SSRN: <http://ssrn.com/abstract=1000533>
- Png, I., I.H. Hann, K.L. Hui, and T.S. Lee, 2008. "Consumer Privacy and Marketing Avoidance: A Static Model." *Management Science*, 54(6), 1094-1103.
- Posner, R. A., 1978. "The Right of Privacy." *Georgia Law Review*, 12(3), 393-422.
- Posner, R. A., 1981. "The Economics of Privacy." *American Economic Review*, 71(2), 405-409.
- Rose, E., 2005. "Data Users versus Data Subjects: Are Consumers Willing to Pay for Property Rights to Personal Information?" *Proceedings of the 38th Hawaii International Conference on System Sciences*.
- Rubin, P.H. and T. M. Lenard, 2002. *Privacy and the Commercial Use of Personal Information*. The Progress & Freedom Foundation, Washington, DC, USA.
- Schwarz, N., 1999. "Self-reports: How the Questions Shape the Answers." *American Psychologist*, 54(2), 93-105.
- Sheehan, K.B., 1999. "An Investigation of Gender Differences in On-Line Privacy Concerns and Resultant Behaviors." *Journal of Interactive Marketing* 13(4), 24-38.
- Sheehan, K.B., 2002. "Toward a Typology of Internet Users and Online Privacy Concerns." *The Information Society*, 18(1), 21-32.
- Slovic P., 1995. "The Construction of Preference." *American Psychologist*, 50(5), 364-71.
- Solove, D., 2004. *The Digital Person: Technology and Privacy in the Information Age*. New York: New York University Press.

- Spiekermann, S., J. Grossklags, and B. Berendt, 2001. "E-privacy In 2nd Generation E-Commerce: Privacy Preferences Versus Actual Behavior." *Proceedings of the ACM Conference on Electronic Commerce*, 38–47.
- Stigler, G. J., 1980. "An Introduction to Privacy in Economics and Politics." *Journal of Legal Studies*, 9, 623–644.
- Stone, E.F. and D.L. Stone, 1990. "Privacy in Organizations: Theoretical Issues, Research Findings, And Protection Mechanisms." In: *Research in Personnel and Human Resources Management*, K.M. Rowland and G.R. Ferries (eds), Greenwich, CT: JAI Press, Vol. 8.
- Taylor, C.R., 2004. "Consumer Privacy and the Market For Customer Information." *RAND Journal of Economics*, 35(4), 631-650.
- Tang, Z., Y. Hu, M. D. Smith, 2007. "Gaining Trust through Online Privacy Protection: Self Regulation, Mandatory Standards, or Caveat Emptor." *Journal of Management Information Systems*, 24(4), 152-173.
- Tedeschi, B., 2002. "Everybody Talks About Online Privacy, But Few Do Anything About it." *New York Times*, June 3.
- Thaler, R. 1980. "Toward A Positive Theory of Consumer Choice." *Journal of Economic Behavior & Organization*, 1(1), 39-60.
- Tsai, J., S. Egelman, L. Cranor, and A. Acquisti, 2011. "The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study." *Information Systems Research*, 22, 254-268.
- Tversky A. and D. Kahneman, 1974. "The Framing of Decisions and the Psychology of Choice." *Science*, 211(4481), 453-8.
- Tversky A., P. Slovic, and D. Kahneman, 1990. "The Causes of Preference Reversal." *American Economic Review*, 80(1), 204-17.
- 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.
- Varian, H.R., F. Wallenberg, and G. Wroch, 2005. "The demographics of the do-not-call list." *IEEE Security & Privacy*, 3(1), 34-39.
- Wathieu, L. and A. Friedman, 2007. "An Empirical Approach to Understanding Privacy Valuation." Harvard Business School Working Paper Number: 07-075.

Notes

¹ Acknowledgment note: the authors of this manuscript are Alessandro Acquisti (Associate Professor, Heinz College, Carnegie Mellon University; acquisti@andrew.cmu.edu), Leslie John (Assistant Professor, Harvard Business School, Harvard; ljohn@hbs.edu), and George Loewenstein (Herbert A. Simon Professor of Economics and Psychology, Department of Social and Decision Sciences, Carnegie Mellon University; gl20@andrew.cmu.edu).

² 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) in order to avoid confounds associated with potential previous disclosures of *existing* personal data.

³ Naturally, if a subject's valuation of her personal data were, for instance, 50 cents, 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 to keep 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 case, the proportion of people who choose the trackable card over the untrackable card should also remain the same across conditions.

⁴ Sheehan (1999, 2002) has highlighted age and gender differences in privacy concerns. We do not use a dummy for gender in this regression since, as noted, Experiment 1 focused on a female population.

⁵ In the [WTA/ Δ 2 Control] condition 45.8% of subjects claimed they would keep the \$10 card, compared to [WTA/ Δ 2], where 60.0% said they would keep their card. Although this suggests that a subtle difference in wording (i.e. cannot be tracked vs. will not be tracked) may have mattered, the difference between the conditions was not statistically significant (Pearson $\chi^2(1) = 1.97$, $p = 0.16$). To continue the analysis of the experiment as a 2x2 factorial design, the [WTA/ Δ 2 Control] condition is excluded from the statistical analyses that follow.

⁶ While the response options presented to the subjects were, necessarily, not evenly spaced, subjects nevertheless had to make discrete choices for each interval. Hence, such spacing cannot explain, alone, the modal points of the distribution, and it does not affect the statistical tests which we present further in the text and that we used to test for normality and unimodality.

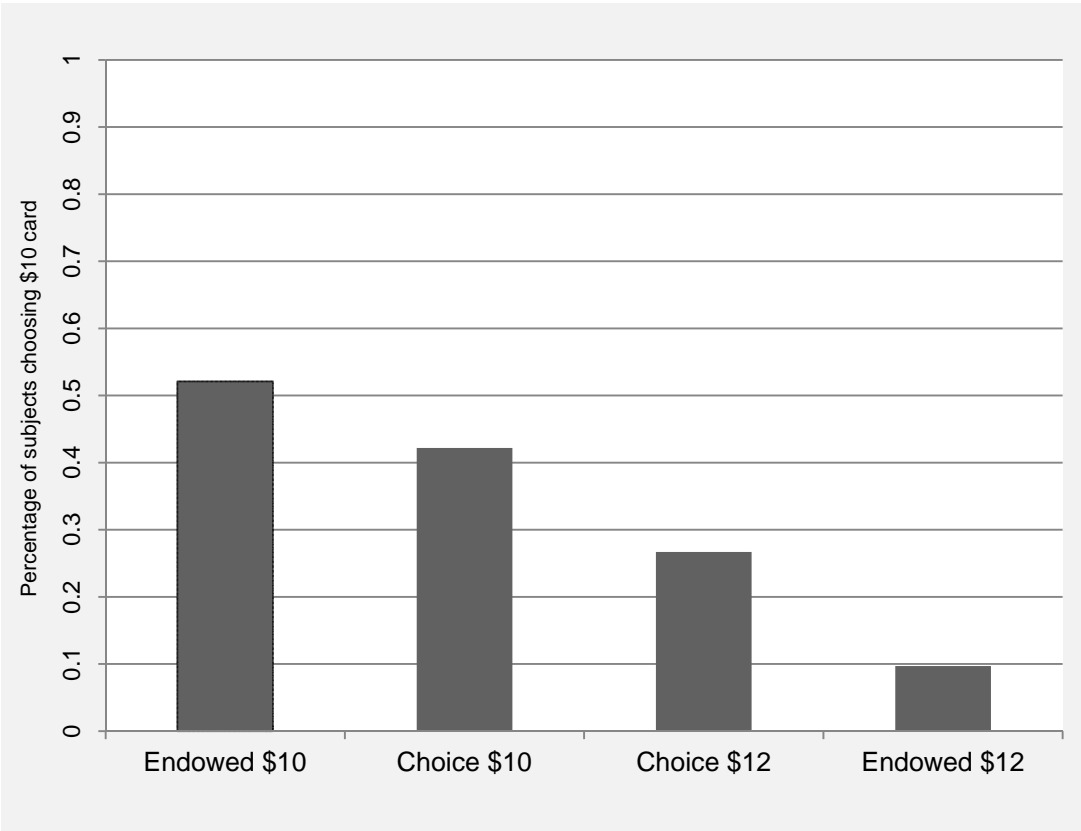


Figure 1

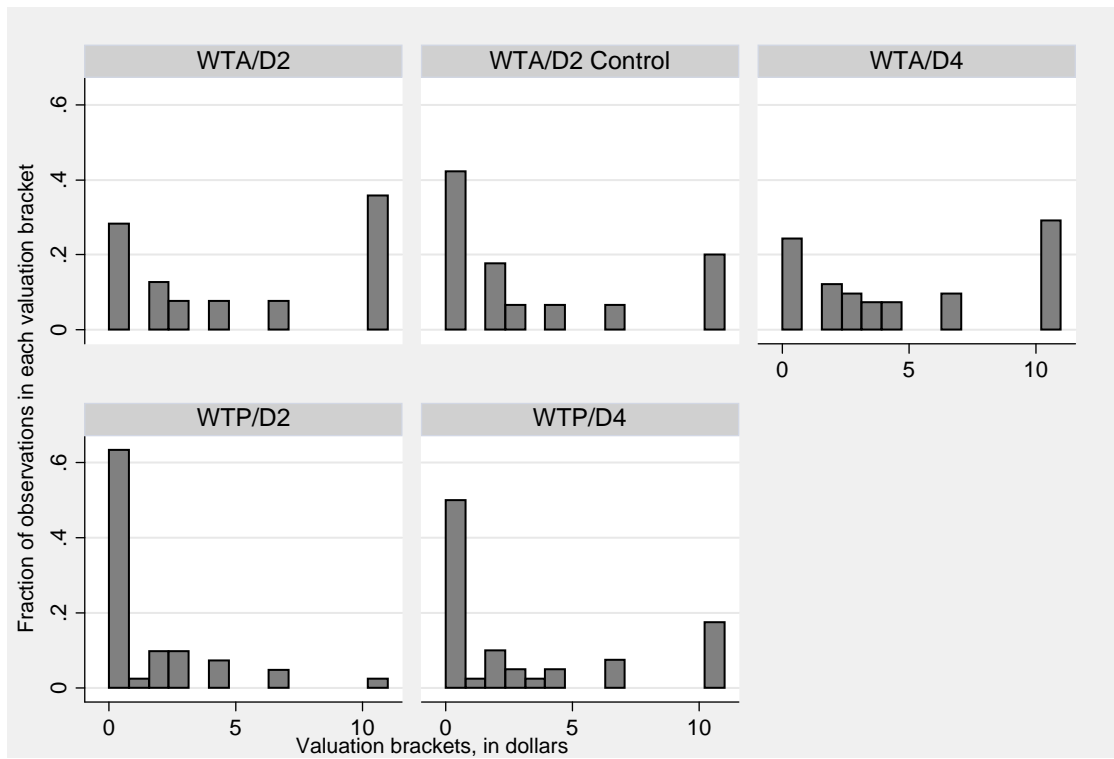


Figure 2

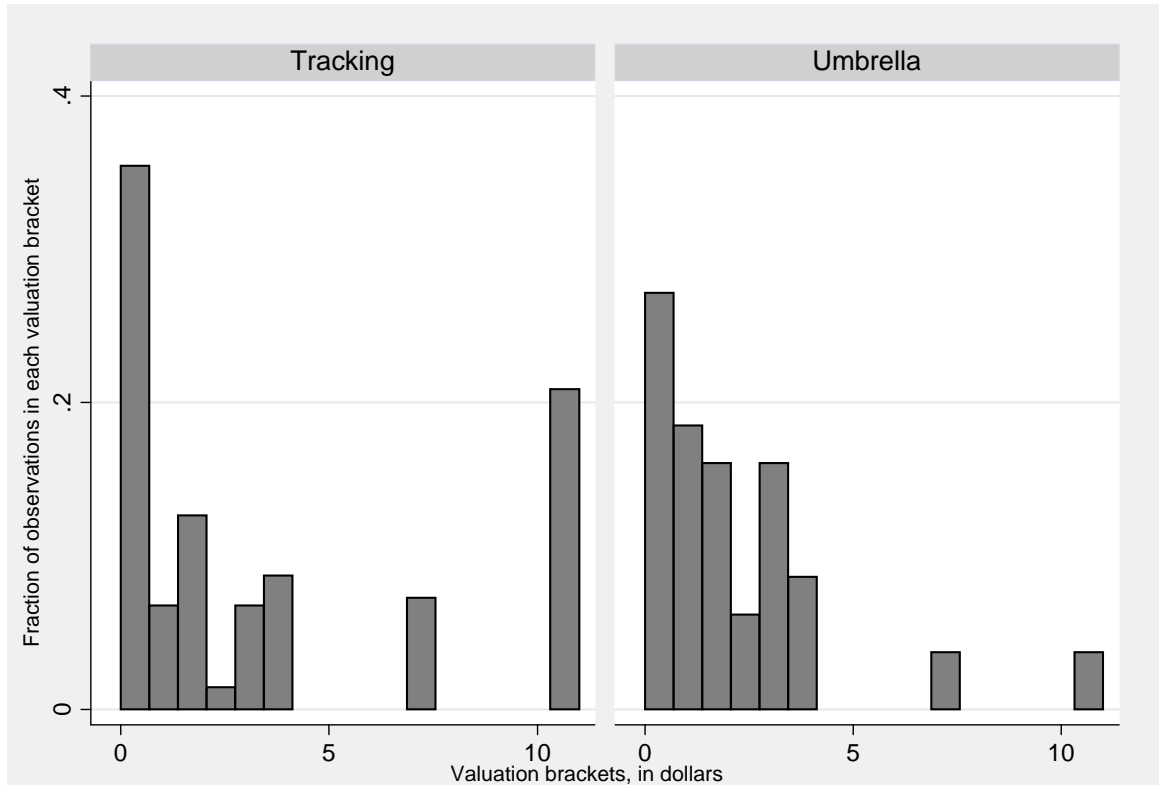


Figure 3

Figure Legends

Figure 1: Gift card selection in Experiment 1. The Y axis represents the percentage of subjects who chose the \$10 anonymous card across four experimental conditions, represented along the X axis.

Figure 2 - Distribution of valuations of purchase data protection in Experiment 2. In each quadrant, the Y axis represents the fraction of observations in each valuation bracket, and the X axis represents the lower boundary (in dollar terms) of each valuation bracket, from \$0 to \$11.

Figure 3 – Comparison of distributions of valuations in Experiment 2 (trading privacy from tracking for money) and in one Falsification test (trading an umbrella for money). In each quadrant, the Y axis represents the fraction of observations in each valuation bracket, and the X axis represents the lower boundary (in dollar terms) of each valuation bracket, from \$0 to \$11.

Table 1 - Probit regression, Experiment 1. The dependent variable represents the card selection (0=\$10 anonymous card, 1= \$12 identified card)

	<i>Endowed conditions</i>	<i>Choice Conditions</i>
Constant	2.4379** (0.4880)	1.1130** (0.3608)
Age	-0.0304** (0.0104)	-.0102 (0.0082)
\$10Card	-1.4400** (0.2917)	-0.6210 ⁺ (0.2417)
	<i>N</i> = 123	<i>N</i> = 128
	<i>Prob</i> > <i>chi2</i> (3) = 0.0000	<i>Prob</i> > <i>chi2</i> (3) = 0.0180
	<i>Pseudo R</i> ² = 0.23	<i>Pseudo R</i> ² = 0.05

Notes: +P < .10; * P < .05; ** P < .01. Standard errors in parentheses.

Table 2 - Probit regression, Experiment 2. The dependent variable represents the card selection (0=\$10 untrackable card, 1= \$12 or \$14 trackable card)

Constant	0.9853** (0.3222)	0.9404** (0.3453)
Age	-0.0185** (.0065)	-0.0181** (0.0066)
Gender	-0.0235 (0.1962)	0.0115 (0.1990)
WTA	-0.6093** (0.1942)	-0.5360 ⁺ (0.2817)
$\Delta 2$	0.1105 (0.1954)	0.1844 (0.2844)
WTA * $\Delta 2$		-0.1420 (0.3972)
	<i>N</i> = 179	<i>N</i> = 179
	<i>Prob</i> > <i>chi</i> 2(4) = 0.0008	<i>Prob</i> > <i>chi</i> 2(4) = 0.002
	<i>Pseudo R</i> ² = 0.08	<i>Pseudo R</i> ² = 0.08

Notes: +P < .10; * P < .05; ** P < .01. Standard errors in parentheses.