

What is privacy worth?

Alessandro Acquisti

Leslie John

George Loewenstein

Under review – Please contact the authors for the most recent version.

A previous version of this manuscript was selected by the Future of Privacy Forum in its “Privacy Papers for Policy Makers” competition, and was presented at the Twenty First Workshop on Information Systems and Economics (WISE).

Abstract

We investigate individual privacy valuations in a series of experiments informed by theories from behavioral economics and decision research. We find evidence of order and endowment effects and non-normal distributions of valuations. In particular, we find that individuals assign markedly different values to the privacy of their data depending on the order in which they consider different offers for that data, or whether they consider the amount of money they would accept to disclose otherwise private information, or the amount of money they would pay to protect otherwise public information. We find the gap between such values to be larger than what usually estimated in comparable studies of other private goods. In addition, we find evidence that privacy valuations are not normally or uniformly distributed, but U-shaped, and clustered around extreme, focal values. These results paint a more nuanced and detailed picture of privacy valuations than the one currently in the literature. They suggest that the value of privacy, while not entirely arbitrary, is highly malleable and sensitive to non-normative factors.

1. INTRODUCTION

Understanding the value that individuals assign to the protection of their personal data is of great importance to policy makers, businesses, and researchers. It is 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. It is important to businesses because, by estimating how much consumers value the protection of their personal data, managers try to predict which privacy-enhancing initiatives may become sources of competitive advantage, and which intrusive initiatives may trigger consumers' adverse reactions. Finally, it is important to researchers, who are interested in measuring the value that individuals assign to privacy, so as to better understand the drivers of information disclosure and information protection.

In recent years, there has been no shortage of empirical studies attempting to quantify individual privacy valuations in diverse contexts (such as online privacy: Hann *et al.* [2007]; location data privacy: 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 growing 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 recently claimed (Gonsalves [2010])?

In this manuscript we challenge the view that true privacy valuations can be precisely estimated, and argue that revealed preferences argument do not necessarily support the conclusion that people, on average, do not care for privacy. Implicit in the current empirical literature on privacy - as published in

information system, economics, marketing, and computer science journals - is the premise that individuals have stable, coherent, and therefore quantifiable valuations of the protection of their data. There are reasons to believe, however, that, consumers' preferences for privacy may not be stable, or even internally coherent. 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 they are quantifiable because they lead to some measurable economic cost, the magnitude and risk of incurring this cost is often uncertain and difficult to quantify (Acquisti [2004]) It would therefore be reasonable to conjecture that valuations of privacy will be strongly 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]).

In this paper, we use a series of experiments to understand privacy valuations (and privacy decision making) through the lenses of behavioral economics and decision research. We find that privacy valuations are, in fact, internally inconsistent and highly vulnerable to subtle, non-normative influences. Specifically, we show the impact of both order effects (Schwarz [1999]) and endowment effects (Kahneman and Tversky [1979]) on privacy valuations, and we highlight a dramatic gap between subjects' "willingness to pay" to protect the privacy of their data and their "willingness to accept" money in order to give up privacy protection. Moreover, we show that this gap is significantly larger than those observed in similar studies on ordinary private goods. These results stand in contrast to other estimates of privacy valuations proposed in the literature, as well as to the view that an economic revealed preferences argument (such as consumers' eagerness to disseminate personal information to friends and strangers alike, or their relative disinterest in freely available protective technologies, such as Tor or PGP) should support the conclusion that consumers do not care for privacy. In one of our experiments, 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 addition, our results offer a more nuanced and detailed understanding of privacy valuations than what currently available in the literature. Numerous studies have tried to pin-point privacy valuations. However, such studies have only focused on mean valuations (or, at best, minimum and maximum values), without casting a light on the actual underlying distribution of those valuations. In this paper, we show that privacy valuations are not normally or uniformly distributed, but U-shaped, clustered 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.

The policy, managerial, and research implications of these findings are significant, because, as noted, privacy valuations matter to policy makers, businesses and researchers. Perhaps even more importantly, by showing that non-normative factors can significantly affect privacy decision making, our findings raise doubts about individuals' ability to optimally negotiate their privacy preferences in today's complex information environment.

2. BACKGROUND

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, since the late 1970s, with the contributions of Hirshleifer (1980) and Chicago School scholars such as Posner (1978) 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, 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

As for all goods and services, there are two common methods of assessing the monetary value that people place on privacy. Applied to privacy, *willingness to accept* (WTA; also known as compensating variation) asks how much an individual would need to be compensated to permit a decrease in privacy. Willingness to pay (WTP; equivalent variation) asks how much an individual would pay to experience an increment in privacy protection. Most empirical efforts to pinpoint individuals' monetary valuations of privacy have focused, either explicitly or implicitly (via the authors' unstated assumptions), on WTA: the willingness to accept payment in exchange for disclosing otherwise private information. Huberman *et al.* (2006) used a second-price auction to study the amount of money individuals would require to their weight or height to others. Wathieu and Friedman (2005) showed that survey participants were comfortable with an institution's sharing 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 - which is somewhat surprising, given the usual finding that WTA tends to produce valuations

that are substantially higher than those produced by WTP. 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 of WTP – that is, privacy valuations in which consumers are, instead, asked to consider paying (or giving up) money to protect their data – are 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. Acquisti and Grossklags (2005) reported that, among survey respondents who believed that technology should be used to protect privacy, the majority of them had decided not to incur the costs of using protective technologies such as encryption or shredders. On the other hand, Tsai *et al.* (2009) 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.

2.1 Privacy valuations

Challenging the assumption that privacy preferences are coherent and consistent (implicit in attempts to measure *the* value placed on privacy), anecdotal evidence suggests that individuals' privacy preferences might not be stable. For example, in surveys, American consumers tend to claim that they are very concerned about their privacy (e.g. Harris Interactive [2001]). Yet, as noted, empirical studies suggest that even self-professed privacy-conscious subjects are willing to provide highly personal information for relatively small rewards, fueling a debate on the existence and nature of a discrepancy between privacy attitudes and privacy behavior (see Acquisti [2004], Wathieu and Friedman [2005], Norberg *et al.* [2007], and Rifon *et al.* [2008]). Despite these hints, no published empirical study has highlighted the impact of non-normative factors on privacy valuations, nor has applied lessons from behavioral economics and

decision research to the measurement of privacy valuations and concerns. For instance, none of the abovementioned manuscripts has explicitly contrasted individuals' willingness to pay to protect data to their willingness to accept money to reveal the same data. 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 analysis cannot distinguish between how much people will pay to protect their data, and how much they will accept to give their data away. Therefore, it cannot determine conclusively 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.

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 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 that data is made available to, or can be managed by, the consumers who generated it (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 instance, the usage of a grocery loyalty card (which creates a potentially sensitive history of purchases at a given store, in exchange for a monetary discount) is attached to the completion of grocery shopping, making it hard to separate consumers' valuations of privacy from their valuation of discounts and purchased goods.

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 how robust are the predictions based on those assumptions. Our research shows that privacy valuations are, indeed, 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, if such a scenario were true, would be the “true” value of the privacy of our data? Both cannot simultaneously reflect our true preferences.

2.2 Why privacy valuations may be malleable and inconsistent

The notion that preference for privacy may not only be context-dependent, but malleable and uncertain, 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 [2009]) on privacy decisions, and how those decisions deviate from the abstractly rational path predicated by neoclassical economic theory.

In this paper, we investigate such deviations as they apply to the privacy domain. First, contrary to traditional economic theory, we consider the significant and robust discrepancy that experimental studies have uncovered between the *maximum* price a person would be willing to pay to acquire a good she did not own (WTP), and the *lowest* price she would be willing to accept to part with the same good if

she initially owned it (WTA). The effect has been replicated time and again (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]) - despite valiant attempts at eliminating it (Plott and Zeiler [2005]), which we further discuss below. Although various explanations have been proposed to explain the discrepancy (Hanemann [1991]; Hoehn and Randall [1987]), by far the best supported, account of the WTP/WTA discrepancy are the endowment effect and loss aversion - the differential treatment of gains and losses (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. Such distinction is crucial for understanding privacy decision making, because decisions involving privacy 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, hides one’s online behavior, but incurs the user the cost of slower downloads. 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 (including PC utilities and productivity tools, digital media applications, and games and entertainment services) in order to monitor their Internet behavior.¹

Second, we consider that consumers’ decisions are affected by the order in which offers are presented (Brookshire *et al.* [1981], Schwarz [1999]). For example, prior research on consumer choice has found a position effect such that the right-most item of a product array (which is what people tend to look at first) is preferred (Nisbett and Wilson [1977]). Recently, Liu and Aaker (2008) found that when consumers are first asked whether they would donate their time to charity, followed by how much money

¹ See http://www.comscore.com/About_comScore/Methodology/Recruitment, accessed on September 26, 2009.

they would donate to charity, monetary donations are increased, relative to when these questions are posed in the reversed order. Applied to privacy, this anomaly suggests that consumers' valuations of their privacy will depend on the order in which they are asked to reveal information of varying degrees of privacy-sensitivity.

Third, we investigate whether the significant amount of uncertainty and ambiguity associated with privacy trade-offs, coupled with the malleability of privacy concerns, may produce unusual distributions of privacy valuations. John, Acquisti, and Loewenstein (2009) have advanced, tested, and found support for the thesis that people don't ordinarily think of privacy, or, thus, take it into account, in the absence of environmental cues that trigger concern. This leads to some surprising effects – for example, that 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.

2.3 Theory and 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 an online purchase completed via an ordinary credit card, while p^+ could represent the condition where, thanks to some anonymous payment technology – see Chaum [1983]) the consumer's online purchases remain private information, and neither the merchant nor, say, the consumer's credit card company, can link the consumer to her purchases. 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). This motivates why we gave subjects in our experiments binary choices of this type (Section 3.1), in addition to actually eliciting exact WTP and WTA values (Section 3.2). Such binary choices are, in fact, much more characteristic of 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, faced with the option of obtaining money in exchange for reduced privacy (WTA), will*

reject it, is larger than the fraction of consumers who, faced with an economically equivalent option of paying for increased privacy (WTP), will accept it.

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 completely arbitrary, are malleable to non-normative factors, and can be, in fact, internally inconsistent, in that the same cash-for-data offer may be accepted or rejected for non-normative reasons.

A related hypothesis would suggest that, if privacy valuations are so malleable, the order in which privacy-sensitive choices are presented may also play a role in the decision by consumers (Schwarz [1999]). Specifically, presenting a privacy enhanced option before the less privacy enhancing one may be interpreted as a signal that the former is inherently more valuable:

(Hypothesis 2) Order effects in privacy valuations: *The fraction of consumers who will choose a privacy enhanced card, faced with the choice between gift cards of different monetary values and privacy features, is larger when the privacy enhanced card is presented before its alternative.*

Finally, if privacy valuations and concerns are malleable and privacy trade-offs uncertain, they may also be influenced by visceral preferences as much as by market-based reasoning. For instance, Westin (1991) famously identified three clusters of consumers as “unconcerned” (those who claim not to care for privacy), “fundamentalists” (those for whom privacy is a fundamental right), and “pragmatist” (those in between the previous two categories). Based on this, we made the following conjecture:

(Conjecture 1) Non-normality of valuations: Unlike most ordinary private goods, the distribution of privacy valuations tends not to be normal or uniform.

3. THE EXPERIMENTS

To test our hypotheses, we ran a series of experiments informed by a common design: subjects were asked to choose between gift cards that varied with respect to their privacy features and monetary value. All experiments investigated subjects' willingness to keep or exchange gift cards as a function of their initial endowment or as a function of the order in which choices were presented. Experiment 1 tested Hypotheses 1 and 2 in the field, with real gift cards. Experiment 2, which consisted of a hypothetical survey, extended Experiment 1's results and allowed us to estimate a distribution of privacy valuations and investigate Conjecture 1.

3.1 Experiment 1: Order and endowment 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. We used \$12 (trackable) and \$10 (untrackable) cards.

Shoppers at a Pittsburgh shopping mall were stopped by research assistants who were blind to the hypotheses of the study. Subjects were offered gift cards in exchange for participating in a survey. The survey was a decoy, intended to create a credible reason for giving the subjects a reward (the gift card), and was identical across all conditions. Subjects across all conditions were asked 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 study was a five condition between-subjects design. There were two "endowed" conditions, two "choice" conditions, and a control condition. In the endowed conditions, subjects were either endowed with the \$10 anonymous card or the \$12 identified card before being offered 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 the 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 other words, 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 and thus is clearly the dominant choice. This condition was included to ensure that people understood and paid attention to the task. We summarize the five 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 on the page) and an identified \$12 card (appearing second on the page)*
4. *[\$12 Choice] Choose between an identified \$12 card (appearing first on the page) and an anonymous \$10 card (appearing second on the page)*
5. *[Control condition] Choose between an identified \$10 card (appearing first on the page) and an anonymous \$12 card (appearing second on the page)*

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: choosing between a \$10 anonymous card or a \$12 identified card. However, the gift card endowment in two of the conditions generated a subtly different framing of the choice faced by the participants: for subjects 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. If privacy preferences were stable and consistent, 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.²

3.1.1 Procedure

Experiment 1 took place on three weekend days at a Pittsburgh shopping mall. Female research assistants were located 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 meant 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. Coupon redemption and subsequent gift card distribution always took place as subjects exited the store. The two endowed conditions and the \$10 choice condition were run during the first weekend. The \$12 choice and the control conditions were run the following weekend. There were five different coupons, each corresponding to a study condition (see Appendix A). To avoid making the different conditions salient, the experimenters distributed coupons for a single condition at a time, rotating the coupon type (and therefore the experimental condition) every hour. Our results (and in particular the card selection) were *not* affected by the time of day when the experiment was conducted, the store in front of which subjects were recruited, or whether the unrelated survey was completed before or after entering the store.

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 in order to increase the potency of the privacy-laden gift card value manipulation. Because the \$10 and \$12 gift cards looked identical,

² 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 participants 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.

they were each labeled with a small, removable sticker that said either “\$10” or “\$12”, as appropriate. The stickers also enabled each card to be tracked. Each card had a unique card number and security code which were recorded in advance. Each card number was then assigned a unique 3-digit number which was written on the sticky side of the label stickers. Once a subject had selected a gift card, the sticker was removed and stuck onto the receipt. Thus, the sticker validated the receipt amount, while also enabling us to track every card’s purchases (subjects could not notice this, since the information was printed on the reverse, sticky side of the sticker).

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. Note that, up until now, 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.

Then, 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. In the \$12 endowed, identified condition, the experimenter recorded the card’s number and security code on the receipt that also contained the person’s name. Next, the experimenter gave the subject a second sheet of paper describing the privacy features of the other, \$10 [\$12] 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. If so, she placed her initial card back into the bin from which she had drawn it, and chose a new one from the other bin. For those in the \$10 endowed condition who exchanged their card, the experimenter recorded the card number and security code of the new, \$12 identified 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. The experimenter recorded the card number and security code for those who chose the \$12 identified card. Once the subject had made her card choice, the

experimenter peeled off the sticker label (also containing the link to the card's number on the sticky side) and stuck it on the receipt. The subject then signed to indicate that she had indeed received the gift card in the value indicated on the sticker. Subjects were then asked to provide their email address.

Note that, across all conditions, subjects had the same amount of time to reflect on how to use their respective cards *in the future*. Specifically, all subjects could mentally compare choosing the trackable card in order to purchase non-sensitive items, versus choosing the anonymous card in order 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=35; M and Median income= \$40,001-\$50,000/year, Mode= \$0-\$10,000; 83.6% Caucasian, 8.5% African American; all not significant between conditions). The majority (92.3%) of subjects returned to the experimenter upon exiting the store 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%) ($\chi^2(1) = 5.14$, $p = 0.023$). However, the likelihood of completing the survey upon entry versus exit did not differ between conditions ($\chi^2(4) = 3.71$, $p = 0.447$), nor did redemption rates were ($\chi^2(4) = 2.35$, $p = 0.673$).

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 we present below.

Overall, most subjects (65.9%) chose the \$12 identified card; however, gift card choice was significantly different across the experimental conditions ($\chi^2(3) = 30.61$, $p < 0.0005$). 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 finally 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 with which they had been endowed; however, 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 ($\chi^2(1) = 27.24, p < 0.0005$). 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 2 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 – in which only order in which the cards were described was varied – 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).³ 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 ($p < 0.0005$): 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 ($p = 0.10$), 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.

³ 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.




Figure 1 - Percentage of subjects who chose, chose, kept, or switched to the \$10 anonymous card in Experiment 2 (vertical axis).

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 $\chi^2(1) = 1.28$, $p = 0.258$), did not have an impact on their likelihood of using the card.

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

| | <u>Endowed conditions</u> | <u>Choice Conditions</u> |
|-----------------|---------------------------|---------------------------|
| 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) |
| | $N = 123$ | $N = 128$ |
| | $Prob > chi2(3) = 0.0000$ | $Prob > chi2(3) = 0.0180$ |
| | $Pseudo R^2 = 0.23$ | $Pseudo R^2 = 0.05$ |

Notes: * $p < .1$, ** $p < .05$, *** $p < 0.01$. Standard errors in parentheses.

We also investigated whether subjects used their cards at different types of stores, depending on card identifiability. On the one hand, subjects who had chosen anonymous cards might be more likely to use them at sensitive stores; on the other hand, it could be that those who are not privacy conscious are both more likely to choose a trackable card *and* to shop at sensitive stores. The latter hypothesis received some support. Stores were classified as potentially privacy sensitive (e.g. lingerie stores such as “Victoria’s Secret”) or not (cafes, drugstores, supermarkets). We found some anecdotal evidence of differences: 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: subjects could use their cards literally at any online or offline store of their choice – making it impossible to design a controlled experiment of their store selection.

Discussion. In Experiment 1, subjects chose different gift cards depending on the framing of the choice, and therefore implicitly assigned dramatically different values to the privacy of their data. Choices in the two endowed conditions were different from the choice conditions, and the choice conditions differed between themselves based on which option was presented first. These patterns stand in contrast to results in the literature looking for and identifying an objective *true* valuation of privacy to be captured. Consider the following: More than half of subjects in the anonymous \$10 endowed condition rejected an offer of \$2 to reveal their future purchase data – in other words, 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. Their WTA was therefore larger than (or at best equal to) \$2 (apparently, these subjects felt “endowed” with the protection of their information). By contrast, fewer than 10% of subjects in the identified \$12 endowed condition gave up \$2 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. This implies that *five times more subjects* chose privacy in one condition over the other, even though they all faced exactly the same choice.

Exploiting a number of simplifying assumptions, we can then 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 extreme values (such as, in our study, \$0 and \$2). This is not an implausible assumption, given the results from Experiment 2 which we present in the following section (also, choosing values higher than \$2 would not weaken, but strengthen, the point we are about to make). 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). The gap between privacy WTP and WTA is notable because, as discussed, 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, and the reliance on “contextual matching” arguments (Payne *et al* [1999]) to resolve problems of preferences uncertainty quite unsatisfactory.

Our results call for caution in the interpretation of market based experiments and analyses of privacy valuations that do not explicitly control for how privacy vs. cash trade-offs are framed. In particular, they call into question the reliance on “revealed preferences” arguments to conclude that consumers do not care for privacy (Rubin and Lenard [2002]), simply because few users are taking advantage of protective solutions available online. 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 the one 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]). Our results therefore imply 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.

Alternative explanations. Some authors have altogether dismissed the existence of an endowment effect (see Plott and Zeiler [2005]), criticizing experiments in this area for their ostensible lack of incentive-compatible mechanisms or the absence of opportunities for subjects to “learn” during the experiment (due to insufficient repetitions of the experimental procedure). Interestingly, the conditions of our experiment, which deviate from the traditional means through which endowment effects are measured, offer a novel counter argument in that debate. Incentive-compatibility was ensured by design (subjects had to, literally, trade private data for cash). The one-off selection of a real gift card (as opposed to a repeated sequence of valuations, as suggested by Plott and Zeiler [2005]), is a more realistic depiction of the privacy decisions made in daily life: first, privacy trade-offs, albeit very frequent, are unique, because of the ever-changing contextual conditions in which personal information is exchanged; and second, as discussed, consumers are rarely able to negotiate the price of their data: they are typically given binary choices, including take-it-or-leave it options. Furthermore, our design markedly decreased the risk that subjects, faced by a menu of choices, could have inferred the study’s goal, thus compromising its results. (The one-off card selection did limit the information we could extract about individual valuations, compared to direct elicitation studies in which multiple choices are presented; however, Experiment 2, below, removes even this possible drawback by measuring *ranges* of privacy valuations.)

Our results cannot be explained by other effects that are sometimes cited as causes for inconsistent valuations of goods. For instance, *status quo* bias (Samuelson and Zeckhauser [1998]: subjects tend to stick to the option they are assigned to), default effects (the initial endowment may have been interpreted as the option that most people take), trade-off avoidance (Luce [1998]: experimental subjects may have not liked the idea of trading-off their cards), or social norms attached to gift giving (subjects may have not wanted to offend the researchers by rejecting the cards they had been offered), *cannot* account for why subjects would have deviated from such norms, perceived defaults, *status quo*, or trade-offs, in such markedly different proportions across the conditions.

Finally, the value of a private card *relative* to the monetary amount with which subjects were initially endowed is an additional factor that may have influenced card choices. Behavioral marketing and economic research have shown that individuals tend to value goods in relative rather than absolute terms (Kahneman and Tversky [1979], Chen *et al.* [1998]). For subjects in the \$10 endowed condition, the opportunity cost of protecting privacy by not switching to a trackable card (\$2) represented a hefty 20% of their initial endowment – and yet, more than half of those subjects chose to pay that cost. In contrast, for subjects in the \$12 endowed condition, the cost of protecting their privacy (again, \$2) amounted to less than 17% of their initial endowment. However, fewer than 10 percent of those subjects chose to take that cost. These comparisons show that our results are robust, and even more significant, when considering relative estimations of the value of privacy.

3.2 Experiment 2: The distribution of privacy valuations

Experiment 2 was a hypothetical survey in two parts. 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 swap that card for a card of different type and value. This first part was similar to Experiment 1, but differed from it in that subjects were asked to choose between \$10 cards with privacy, and \$12 *or* \$14 cards without privacy: a first goal of Experiment 2 was, in fact, to test whether the WTP/WTa dichotomy uncovered by Experiment 1 would extend to cases where the differential cash value in the card was larger than \$2. More importantly, the second part of Experiment 2 allowed us to estimate subjects' distributions of privacy valuations: after stating which card they would keep, 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 two by two 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. Therefore, 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).

Experiment 2 was run at cafeterias in hospitals in a urban center. Subjects were recruited on site, and were offered chocolate bars to complete the hypothetical questionnaire. Two hundred and forty subjects participated in the study (46.2% female; M age=83, sd =15, Median=35, range=19-83; 75.0% Caucasian) and were randomly assigned to the five experimental conditions (50 subjects participated in condition [WTA/Δ2], 45 in condition [WTA/Δ4], 51 in condition [WTP/Δ2], 44 in condition [WTP/Δ4], and 50 in the [WTA/Δ2 Control] condition). Except for a slight overrepresentation of females in Condition [WTA/Δ2],⁴ there were no other significant demographic differences between conditions.

The first page of the questionnaire stated that the gift cards came in two forms: trackable or untrackable (Appendix B). 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

⁴ We included gender and age in the regression analyses presented below. However, we did not observe any gender effect on card’s choice.

asked whether they would like to keep card they were initially offered, or exchange it for the other card. Answers to those questions allowed us to test whether we could find evidence of an endowment effect when cards values differed by \$2 and by \$4. After answering the question on the first page of the questionnaire, subjects were instructed to turn the page and answer the follow-up questions that allowed us to estimate the subject's distribution of privacy valuations. On the last page, subjects answered demographic questions.

3.2.2 Results

In the conditions in which we asked subjects to choose between a \$10 anonymous card and \$12 trackable card (conditions [WTA/ $\Delta 2$] and [WTP/ $\Delta 2$]), 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/ $\Delta 4$] and [WTP/ $\Delta 4$]): 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 ($\chi^2(1) = 2.95$, $p = 0.086$).

⁵ In the [WTA/ $\Delta 2$ Control] condition 45.8% of subjects claimed they would keep the \$10 card, compared to [WTA/ $\Delta 2$], 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/ $\Delta 2$ Control] condition is excluded from the statistical analyses that follow.

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>chi2</i> (4) = 0.0008 | <i>Prob</i> > <i>chi2</i> (4) = 0.002 |
| | <i>Pseudo R</i> ² = 0.08 | <i>Pseudo R</i> ² = 0.08 |

Notes: * $p < .1$, ** $p < .05$, *** $p < 0.01$. Standard errors in parentheses.

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 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 43.2% in the latter case claimed they would switch; Pearson $\chi^2(1) = 0.91$, $p = 0.339$). These results are surprising, in that they point to an almost binary attitude towards privacy that is powerfully affected by WTA and WTP, but not by monetary differences. This could be possible if, in the context of the experiment, privacy valuations did not vary much in the [\$2-\$4] interval. For instance, some individuals may value such 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 not value such protection at all (less than \$2, so being offered \$2 or \$4 would not make a difference to them either); but very few individuals value the privacy of that purchase data *exactly* \$x, with $2 < x < 4$ – hence the lack of difference in selection patterns in the \$10 versus \$14 conditions over the \$10 versus \$12 conditions in Experiment 2. This interpretation of our findings would be compatible with the conjecture that privacy valuations are not uniformly or even normally distributed, but 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 such conjecture.

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 B). Subjects in the WTA conditions who chose to keep an untrackable \$10 card were asked: "Would you have also kept the card you were originally given if it had been a [\$9.75, 9.50, 9.25, 9, 8, 5, 1] card that will not be tracked?" Subjects in the WTA Conditions who instead chose to exchange a \$10 card for a \$12 card were asked: "Would you have also exchanged the card you were originally given for a [\$11.75, 11.50, 11.25, 11, 10.75, 10.50, 10.25] card that will be tracked?" Subjects in the WTP Conditions who

chose to keep the \$12 trackable card were asked: “Would you have also kept the card you were originally given if it had been a \$[11.75, 11.50, 11.25, 11, 10.75, 10.50, 10.25] card that will be tracked?” Subjects in WTA Conditions who chose to exchange the \$12 trackable card for a \$10 untrackable card were asked: “Would you have also exchanged the card you were originally given for a \$[9.75, 9.50, 9.25, 9, 8, 5, 1] card that will not be tracked?”⁶

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). They must value the privacy of their transaction data at least \$2. Among them, consider the person who 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).⁷ We then took the lower boundary of each bracket, and constructed the histograms presented in Figure 3 (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).

⁶ Subjects in the Δ4 Conditions answered similar questions, only that the values presented in the follow-up questionnaire were calibrated on the different value of their trackable card; see Appendix B.

⁷ Similarly, consider the subjects who chose to keep a \$12 trackable card (rather than switching to a \$10 untrackable card). We already know that these subjects must value their privacy in that particular context less than \$2. Among them, now consider the person who went on to indicate that he would have also kept the trackable card if it had been worth \$11.50, but not if it had been worth \$11.25. In this case, we would then infer him to have a (self-reported) valuation for privacy of no more than \$1.50, but no less than \$1.25.




Figure 2 - Distribution of point-wise valuations of purchase data protection based on the results of Experiment 2. The vertical axis represents the fraction of observations in each range of valuations. The horizontal axis represents identical value across the quadrants: the lower boundary (in dollar terms) of each valuation bracket, from \$0 to \$11.

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/ Δ 2 condition) are markedly U-shaped (also, consistent with Hypothesis 1, the U-shape is more accentuated in the conditions in which subjects were endowed with the privacy enhanced card).⁸ 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.⁹ Shapiro-Wilk, Shapiro-Francia, and Skewness-Kurtosis tests on the bracket data all strongly rejected the hypothesis of normality

⁸ 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 ($z = 3.67$, $p < 0.0005$) 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 .

⁹ 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.

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 for conditions [WTA/ $\Delta 2$] and [WTA/ $\Delta 4$] and the Control condition ($p < 0.0005$), 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 3 - Distribution of point-wise valuations: comparison between Experiment 2 conditions (trading privacy for money) and the falsification test conditions (trading an umbrella for money).

As a falsification test, we ran a new battery of experiments using the exact same language in Experiment 1's [WTA/ $\Delta 2$] and [WTP/ $\Delta 2$] conditions, but asking 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. In three different online experiments, we tested subjects' reactions to goods whose average eBay price we found to fall 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 over those physical goods using the design of Experiment 2, the bimodality of the distributions disappears. As an example, consider Figure 3, which is based on the "IKEA umbrella" falsification test: 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 ($\text{diptest}_{[\text{WTP/umbrella}]}: p = 0.28$; $\text{diptest}_{[\text{WTA/umbrella}]}: p = 0.10$).

4. IMPLICATIONS

The combined results of Experiments 1 and 2 paint a more nuanced and granular picture of privacy valuations than currently accepted: privacy valuations, while not completely arbitrary, are subject to subtle framing effects and are anchored around extreme focal points. Specifically, valuations are affected by simple order effects; the “price” people assign to protect a piece of information is very different from the price they assign to sell the same piece of information; and valuations are not normally or uniformly distributed, but tend to be U-shaped.

Researchers have correctly noted that privacy is an ambiguous, multi-faceted concept (Solove [2006]). Even when limited to the protection of one's purchase history, there are many, even contradictory, forces which may affect individual valuations of such protection – from the desire to avoid stigma, to the benefits associated with the avoidance of price discrimination in a repeated purchase scenario. Clearly, each of our subjects had their own different motivations for opting for one card versus the other, and therefore different valuations of the protection of their data. This does not contradict our results (thanks to randomization, subjects with different motivations – and valuations – would be similarly distributed across experimental conditions),¹⁰ but opens up a research agenda aimed at further determining how valuations (and their malleability) change as function of the different types of private

¹⁰ Furthermore, as Experiment 1 demonstrated, selecting different monetary values may or may not alter the proportions of subjects choosing either card, but would not invalidate the basic finding of a WTP/WTA dichotomy. Clearly, increasing the monetary gap between trackable and untrackable cards would also increase the proportion of people choosing the higher-valued card. Such a result would not disprove the WTP/WTA dichotomy, but simply demonstrate the existence of boundary valuations beyond which consumers become privacy insensitive.

information being investigated, or of the different mediators and effector (such as the perception of privacy protection as an unalienable right) activated during a study.

Our results stand in contrast to neoclassical economic models of privacy and current empirical results in this area, which typically assume stable and coherent preferences for privacy (see Section 2.1). Hence, a first implication of our results relates to the theoretical economic literature on privacy. In particular, we show that the assumption that trade-offs between privacy and cash are independent of one's initial endowment is untenable. Since the results of economic models are used to influence and direct public policy initiatives, our empirical results carry a practical lesson to guide our efforts as modelers: we need to vet our theories by testing whether their results are robust when the economic agent's valuation for privacy changes with the "direction" of the cash-for-data exchange.

A second implication pertains to the empirical literature on privacy. In their paper on coherent arbitrariness, Ariely, Loewenstein, and Prelec (2003) noted that "demand curves estimated from market data need not reveal true consumer preferences, in any normatively significant sense of the term." Similarly, our findings cast doubt on the ability to infer consumers' exact evaluations of personal privacy from market experiments: what people decide their data is worth depends critically on the context in which they are asked - specifically, how the problem is framed. While this is true of other ordinary private goods, the gap between WTP and WTA is much larger than the average. More importantly, ordinary private goods are routinely and explicitly sold and bought by consumers; hence, for them, some objective prices can be nevertheless determined. Not so for personal data – which is often transacted as a secondary, almost invisible feature of other primary transactions. This state of things currently leads managers and policy makers to rely on published empirical studies that attempt to pinpoint "exact" privacy valuations. We show that these valuations should be interpreted with extreme caution: analyses that do not adequately differentiate between how much an individual would pay, or would accept, for her private data, are going to hide the reality of how malleable and mutable those valuations can be. More granular information about consumers' valuations therefore can aid sounder management and policy making.

Thirdly, this research raises doubts about 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. The finding that non-normative factors powerfully influence individual privacy valuations may signal the appropriateness of policy interventions.

Finally, and perhaps most importantly, this research has policy and managerial implications. Individuals' decisions about their data are sometimes taken as representing their 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 in fact, ask for it and take advantage of opportunities for its protection. However, as we have demonstrated, "revealed preferences" arguments should not, alone, justify the uncritical conclusion that even privacy conscious consumers will always be unlikely to pay for online privacy. If individual decisions regarding privacy are malleable to non-normative factors, then such arguments lose their normative standing.

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: in our experiments, subjects who started from positions of greater privacy protection were five times more likely than other subjects to forego money to preserve that protection. Combined with the difficulty of making the "right" privacy decisions for consumers, such findings suggest that market outcomes alone may not necessarily tell us the final and last words on consumer data protection.

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.
- Acquisti, A. and J. Grossklags, 2005. "Privacy and Rationality in Decision Making," *IEEE Security and Privacy*, 3(1), 26-33.
- Ariely, D., G. Loewenstein, and D. Prelec, 2003. "Coherent Arbitrariness: Stable Demand Curves Without Stable Preferences," *Quarterly Journal of Economics*, 118(1), 73-105.
- Brandimarte, L., A. Acquisti, and G. Loewenstein, 2009. "Privacy Concerns and Information Disclosure: An Illusion of Control Hypothesis," poster, *iConference 2009*.
- 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.
- Chen, S-F. S., K. B. Monroe, and Y.C. Lou, 1998. "The Effects of Framing Price Promotion Messages on Consumers' Perceptions and Purchase Intentions," *Journal of Retailing*, 74(3), 353-372.
- Culnan, M.J. 2005. "Consumer Awareness of Name Removal Procedures: Implication for Direct Marketing," *Journal of Interactive Marketing*, 9, 10-19.
- 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.
- 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.
- 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.
- Harris Interactive, 2001. “Privacy On & Off the Internet: What Consumers Want.” Technical report, http://www.aicpa.org/download/webtrust/priv_rpt_21mar02.pdf.
- 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, 2006. “Valuating Privacy,” *Proceedings of the Workshop on the Economics of Information Security (WEIS '06)*.
- 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 & Used," Consumerunion.org, September 25. Available at http://www.consumersunion.org/pub/core_telecom_and_utilities/006189.html, accessed on January 20, 2010.
- John, L.K., A. Acquisti, and G. Loewenstein, 2009. "The Best of Strangers: Context Dependent Willingness to Divulge Personal Information." Under review. Available at SSRN: <http://ssrn.com/abstract=1430482>.
- 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.
- 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.
- List, J.A. and J.F. Shogren, 1998. "Calibration of the difference between actual and hypothetical valuations in a field experiment," *Journal of Economic Behavior and Organization*, 37(2), 193-205.
- Liu, W. and J. Aaker, 2008. "The Happiness Of Giving: The Time-Ask Effect," *Journal of Consumer Research*, 35(5), 543-557.
- Luce, M.F., 1998. "Choosing to Avoid: Coping with Negatively Emotion-Laden Consumer Decisions," *Journal of Consumer Research*, 24 March, 409-433.
- Nisbett, R. E. and T. Wilsom, 1977. "Telling More Than We Can Know: Verbal Reports On Mental Processes," *Psychological Review*, 84(3), 231-259.

- 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.
- Norberg P.A., D.R. Horne, and D.A. Horne. 2007. "The Privacy Paradox: Personal Information Disclosure Intentions Versus Behaviors" *The Journal of Consumer Affairs*, 41(1), 100-126.
- Payne, J.W., J.R. Bettman, and D.A. Schkade, 1999. "Measuring Constructed Preferences: Towards a Building Code," *Journal of Risk and Uncertainty*, 19(1-3), 243-270.
- 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," 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. "An economic theory of privacy," *Regulation*, May-June, 19-26.
- Rifon, N. J., R.J. LaRose, and M.L. Lewis, 2007. "Resolving the Privacy Paradox: Toward A Social-Cognitive Theory of Consumer Privacy Protection" Mimeo, Michigan State University, <https://www.msu.edu/~wirthch1/privacyparadox07.pdf>
- 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 (HICSS '05)*.
- 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.
- Shostack, A., 2003. "Paying For Privacy: Consumers And Infrastructures." *Proceedings of the Second Annual Workshop on Economics and Information Security (WEIS '03)*, College Park, MD.
- Slovic P., 1995. "The construction of preference," *American Psychologist*, 50(5), 364-71.
- Solove, D. J., 2006. "A Taxonomy of Privacy," *University of Pennsylvania Law Review*, 154(3), 477-560.
- 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.
- 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.
- Taylor, C.R., 2004. "Consumer Privacy And The Market For Customer Information," *RAND Journal of Economics*, 35(4), 631-650.
- Tedeschi, B., 2002. "Everybody Talks About Online Privacy, But Few Do Anything About it." *New York Times*, June 3, Section C, Page 6, Column 1.
- 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, 2009. "The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study," *Information Systems Research*, forthcoming.
- 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.
- Varian, H.R., F. Wallenberg, and G. Woroch, 2005. "The demographics of the do-not-call list," *IEEE Security & Privacy*, 3(1), 34-39.
- Wathieu, L. and A. Friedman, 2005. "An Empirical Approach to Understanding Privacy Valuation," *Proceedings of the Fourth Workshop on the Economics of Information Security (WEIS '05)*, Cambridge, MA, June 2-3.
- Westin, A.F., 1991. "Harris-Equifax Consumer Privacy Survey, 1991." Atlanta, GA: Equifax Inc. 1991.
- Willig, R.D., 1976. "Consumer's Surplus Without Apology," *American Economic Review* 66(4), 589-597.