What is privacy worth?[1]

Alessandro Acquisti, Leslie John, and George Loewenstein
Carnegie Mellon University

PRELIMINARY DRAFT – PLEASE DO NOT DISSEMINATE

Abstract

We investigate individuals' valuations of privacy using field and lab experiments. We find that privacy valuations are inconsistent and highly dependent on subtle framing. Specifically, we find evidence of a dichotomy between "willingness to pay" and "willingness to accept" for privacy: Individuals assign radically different values to the protection of their data, depending on 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. These results suggest that the value of privacy, while not entirely arbitrary, is highly malleable and sensitive to non-normative factors. Therefore, they raise doubts about individuals' ability to optimally negotiate issues of privacy in modern information societies.

[1] Acknowledgments.

What is privacy worth?

Alessandro Acquisti, Leslie John, and George Loewenstein

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 often need 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 knowing how much consumers value the protection of their personal data, firms can predict which privacy enhancing initiatives may become sources of competitive advantage, and which intrusive initiatives may instead trigger consumers' adverse reactions. Finally, it is important to researchers, who have devoted considerable effort to measuring the value that individuals assign to privacy, so as to better understand the drivers of information disclosure and information protection. Is it really possible, however, to measure *the* value that people place on privacy? The premise that privacy valuations can be precisely estimated assumes that individuals have relatively stable and coherent privacy preferences. In this paper, however, we question whether people actually place a consistent, coherent, value on privacy, and hence whether such valuations of privacy can be pinpointed in any meaningful sense. In two experiments, we investigate individuals' trade-offs between money and privacy, and find that privacy valuations are internally inconsistent and highly malleable to subtle, non-normative influences.

In recent years, there has been no shortage of empirical studies attempting to quantify individual valuations of privacy in various contexts (such as online privacy, as in Hann *et al*. [2007], location data privacy, as in Cvrcek *et al.* [2006], or removal from marketers' call lists, as in Varian *et al.* [2005]). Implicit in many of these studies is the assumption that individuals have stable 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 someone 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 risk of experiencing this cost is often delayed and uncertain (Acquisti [2004], Acquisti and Grossklags [2005]), and hence subject to the complexities of discounting for risk and time. Given all of this, it would not be surprising to find valuations of privacy to be subject to many of the effects that have come under the heading of "preference uncertainty" (Slovic [1995]). When preferences are uncertain, research has shown, decision making is likely to be influenced by factors that are difficult to justify on normative bases, such as how alternatives are framed (Tversky and Kahneman [1974]) or preferences are elicited (Tversky, Slovic, and Kahneman [1990]).

In this paper, we focus on one such deviation from normative models of decision making: the typically large discrepancy between individuals' willingness to pay (WTP) money for goods they do not own and their willingness to accept (WTA) money to give up goods they own (Thaler [1980]). We conducted two experiments, a hypothetical choice study and a field study, to investigate the hypothesis that privacy valuations may, in fact, be affected by a similar dichotomy between the willingness to pay to protect, and willingness to accept to disclose, personal data. In both experiments, subjects were offered a choice between gift cards with different characteristics. Some of the cards were identifiable, and therefore trackable; others were anonymous, and therefore untrackable. In one experiment we used hypothetical gift cards; in the other, we used real gift cards. By demonstrating a radical difference 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, our findings suggest that privacy valuations are malleable and sensitive to non-normative factors. Furthermore, we used data from one of the experiments to estimate the distribution of privacy valuations, finding evidence that said valuations may not be uniformly or even normally distributed, but in fact clustered around focal, extreme points.

The policy implications of these findings are significant, because, as noted, privacy valuations matter to policy makers, businesses and researchers. Perhaps even more importantly, however, 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.

The rest of the paper is organized as follows: In Section 2 we discuss the literature on privacy valuations, the literature on the WTP/WTA discrepancy, and the theoretical foundations of our

experimental hypotheses. In Section 3 we present the design of, and results from, the two studies. Section 4 concludes with a discussion of implications and limitations.

2. BACKGROUND

2.1 Privacy valuations

The empirical literature on privacy valuations is closely connected to the theoretical literature on the economics of privacy. 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, 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 Acquisti and Varian [2005], Taylor [2004a, 2004b], Calzolari and Pavan [2006], Tang *et al*. [2007], and Png *et al*. [2008]). At the same time, the management, marketing, and IS literatures were also exploring 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], Culnan and Bies [2003], 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 [2004a]). 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.

Most empirical efforts to pinpoint individuals' monetary valuations of privacy have focused, either explicitly or implicitly through the authors' unstated assumptions, on the willingness to accept payment in exchange for disclosing otherwise private information. For example, Tedeschi (2002) reported on a 2002 Jupiter Research study that found that 82% of online shoppers were

willing to give personal data to new shopping sites in exchange for the chance to win $100. Spiekermann *et al*. (2001) studied subjects' willingness to answer personal questions in order to receive purchase recommendations and discounts. Chellappa and Sin (2005) found evidence of a tradeoff between consumer valuation for personalization and concerns for privacy. Huberman *et al*. (2006) used a second-price auction to study the amount of money individuals would require to reveal personal information (such as their weight or height) to others. Wathieu and Friedman (2005) showed that survey participants were comfortable with an institution's sharing of the personal information, if they had been shown the economic benefits of doing so. Cvrcek *et al.* (2006) studied hypothetical valuations for the release of data tracking a subject's personal location over extended periods of time. 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.

Empirical studies of 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 to actually 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, 63 percent had never used any form of encryption for their data; 44 percent did not use email filtering technologies; and 50 percent did not use shredders for sensitive documents. (These numbers offer indirect measures of individuals' willingness to incur intangible or tangible costs – something akin to paying to protect – to secure their data.) However, Tsai *et al.* (2009) found that once privacy-relevant information was made salient, participants in an experiment paid moderate price premia (of roughly 50 cents) to purchase goods from online merchants with better privacy protection, and Varian *et al.* (2005) and Png (2007), trying to estimate the value that US consumers assign to the protection from telemarketers, found values ranging from a few cents to more than $30.[2]

A number of studies, however, provide hints that individuals' privacy preferences might not be as stable as many researchers have assumed. For example, in surveys, American consumers tend to claim that they are very concerned about their privacy (e.g. Harris Interactive [2001]). Yet,

[2] While in this paper we focus – as most of the economics and IS literatures do – on information privacy, researchers also investigate privacy under its original definition as "the right to be left alone," by Warren and Brandeis (1890). The two latter studies belong to this group.

empirical studies suggest that even self-professed privacy-conscious subjects are willing to provide highly personal information for relatively small rewards (Spiekermann *et al*. [2001], Tedeschi [2002]), fueling a debate on the existence and nature of a discrepancy between privacy attitudes and privacy behavior (see Shostack [2003], Syverson [2003], Acquisti [2004], Wathieu and Friedman [2005], Norberg, Horne, and Horne [2007], and Rifon, LaRose, and Lewis [2008]).[3]

Yet, despite these hints, no published study has explicitly contrasted individuals' willingness to pay to protect data to their willingness to accept money to reveal the same data. In fact, the distinction between the two concepts is usually absent in the literature. For instance, in a seminal contribution on privacy valuations, Hann *et al.* (2007) quantified 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). However, while the conjoint analysis employed in that study does provide a valuable tool to compare individuals' preferences over money and websites features, it cannot distinguish between WTP and WTA, and therefore cannot determine conclusively whether, and how much, individuals will in actuality pay to *protect* their privacy.

2.2 The WTP/WTA discrepancy

Inspired by research in behavioral economics, the notion that preference for privacy may not only be context-dependent, but in fact uncertain, suggests that studies investigating privacy valuations may not tell us much about whether consumers will actually pay to protect their data. First, behavioral economists have highlighted that non-normative factors often affect valuations and decision making under uncertainty (Slovic [1995]). Second, a large body of research has documented a significant and robust discrepancy 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).

[3] Dichotomies between stated privacy attitudes and actual behavior may often be resolved by observing that the framing of a privacy survey is different from the actual trade-offs that consumers faced when making privacy sensitive decisions in real life (in other words, generic privacy intentions are, not surprisingly, poor predictors of specific privacy behavior; see also Fishbein and Ajzen [1975]). However, dichotomies between stated privacy attitudes and specific disclosure behaviors were also uncovered within more narrowly and consistently defined scenarios (for the online social networks case, see Acquisti and Gross [2006]).

Early research in this area suggested that people's stated willingness to pay for public goods displayed properties difficult to explain using standard economic theory, in turn raising skepticism as to whether such measures represent genuine economic preferences (see Hammack and Brown [1974]). Similarly, Kahneman and Knetsch found stated willingness to pay for environmental goods to be wildly insensitive to scope – WTP to clean up one polluted lake does not differ from WTP to pay to clean up all such lakes (see Kahneman [1986]). Building on this research, Kahneman, Knetsch, and Thaler (1991) later demonstrated that, again contrary to traditional economic theory, valuations of products (in their study, mugs) were highly impacted by ownership: subjects who had been endowed with a mug demanded more money to give it up than potential buyers were willing to pay to acquire the same mug; in other words, a substantial WTP-WTA gap, or endowment effect, was observed. The effect has been replicated time and again (Knetsch [1989], Kahneman, Knetsch, and Thaler [1990]), 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]), and has become so well-established that neuroeconomic research has begun to identify its neural underpinnings (Knutson, Wimmer, Rick, Hollon, Prelec, and Loewenstein [2008]).

Various explanations have been proposed in the literature to explain the WTP/WTA discrepancy, including lack of substitutability between goods (see Hanemann [1991]), as well as uncertainty about a good's value and bounded rationality (see Hoehn and Randall [1987], Eisenberger and Weber [1995], and Roth [2005]). By far the most frequent, and best supported, account of the WTP/WTA discrepancy, however, involves the differential treatment of gains and losses – loss aversion (Kahneman and Tversky [1979], Thaler [1980]). Applied to privacy, loss aversion 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.

The distinction between WTP and WTA 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[4] hides one's online behavior, but incurs the user the cost of slower downloads. Analogous to WTA, in other situations we are

[4] See www.torproject.org.

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.[5] The research on the WTP/WTA discrepancy would suggest that we assign different valuations to privacy depending on whether the problem is framed as one of protecting our data or one of disclosing it. 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 we are assigning to the privacy of our data? Both cannot simultaneously reflect our "true" preferences.

2.3 Theory and hypotheses

To motivate our experimental design, 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 instance, p^- represents an online purchase completed via an ordinary credit card; p^+ represents instead the condition where 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).[6] 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], Weber [2003]). 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

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

[6] Technologies for such privacy-enhanced payments actually exist. See, for instance, research in this area based on Chaum (1983).

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, rather than eliciting actual WTP and WTA values (see Sections 3.1 and 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 faced with these two choices. However, 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 that it can be $u(w^-, p^+) > u(w^+, p^-)$ while also, simultaneously, $u(w^+, p^-) > u(w^-, p^+)$, simply depending on how the question is framed. This would suggest the following: 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 her data being 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 have no such endowment and are given the opportunity to pay to obtain it; more broadly, 2) privacy preferences, while not necessarily arbitrary, are malleable to non-normative factors, and can be, in fact, internally inconsistent.

3. THE EXPERIMENTS

We ran two experiments in which subjects were asked to choose between gift cards that varied with respect to their privacy features and monetary value.[7] Both experiments investigated subjects' willingness to keep or exchange gift cards as a function of their initial endowment, and tested Hypothesis 1. Experiment 1 tested the hypothesis by means of a hypothetical questionnaire. Its results also helped us calibrate the values of the cards to be used in Experiment 2, which was a field experiment with real gift cards. An additional purpose of Experiment 1 was to provide us with data points to estimate a distribution of privacy valuations.

3.1 Experiment 1

In Experiment 1, 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.

The first page of the questionnaire stated that the gift cards came in two forms: trackable or untrackable (Appendix A). 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." In the language of Section 2.3, the untrackable card therefore represented a scenario with greater privacy protection than its alternative. Subjects were randomly assigned to experimental conditions that differed by the type of card they were initially offered. Subjects were then asked whether they would like to keep card they were initially offered, or exchange it for the other card.

The experiment was a 2 by 2 between-subjects factorial design. We manipulated a) whether subjects were 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) a subject would be willing to

[7] The definition of the two cards (trackable or identified card versus untrackable or anonymous card) was consistent *within* each experiment, but was slightly different *across* experiments, in order to test the robustness of the findings to different (but equivalent) descriptions of the cards.

pay back to protect her data, and conditions where subjects were assigned an untrackable card as "WTA" since they relate to the question of how much (if anything) a subject would be willing to accept to give away her 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. [WTP/Δ2] Keep $12 card which will be tracked, or exchange for $10 card which cannot be tracked
3. [WTA/Δ4] Keep $10 card which cannot be tracked, or exchange for $14 card which will be tracked
4. [WTP/Δ4] Keep $14 card which will be tracked, or exchange for $10 card which cannot be tracked
5. [WTA/Δ2 Control] Keep $10 card which cannot be tracked, or exchange for $12 card which may be tracked

The fifth condition (a variant of the [WTA/Δ2] condition) was included to test whether subjects may be sensitive to slight changes in the description of the cards. It asked subjects 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).

After answering the question on the first page of the questionnaire, subjects were instructed to turn the page and answer follow-up questions (which we will discuss in Section 3.3). On the last page, subjects answered demographic questions.

Note that all subjects, regardless of the condition to which they had been randomly assigned, in reality had to choose between the very same alternatives: a $10 "untrackable" card or a $12 [$14] "trackable" card. However, for subjects in the WTA conditions, the implicit choice was whether to *sell* one's future purchase data to the researchers for $2 [$4]; for those in the WTP conditions, the implicit choice was whether to *pay* $2 [$4] in order to *avoid* having one's future purchase data made available to the researchers.

Experiment 1 was run at cafeterias in hospitals in the Pittsburgh area in late February 2008. Subjects were recruited on site, and were offered chocolate bars to complete the hypothetical

questionnaire. Two hundred and forty subjects participated in the study and were randomly assigned to the four experimental conditions (50 subjects participated in condition [WTA/Δ2], 51 in condition [WTP/Δ2], 45 in condition [WTA/Δ4], 44 in condition [WTP/Δ2], and 50 in the [WTA/Δ2 Control] condition). Subjects' ages ranged from 19 to 83 (mean: 39; standard deviation: 15, median: 35). Females represented 46.2% of sample, and were slightly overrepresented in Condition 1.[8] The sample was predominantly Caucasian (75.0%). Except for a slight overrepresentation of females in Condition [WTA/Δ2], there were no other significant demographic differences between conditions.

3.1.1 Results

In the [WTA/Δ2 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 χ^2 (1) = 1.97, p = 0.16). To continue the analysis of the experiment as a 2x2 factorial design, the [WTA/Δ2 Control] condition is excluded from the statistical analyses that follow.

In the conditions in which we asked subjects to choose between a $10 anonymous card and $12 trackable card (conditions [WTA/Δ2] and [WTP/Δ2]), we found, as hypothesized, a significant effect of card endowment on card choice. 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 (χ^2 (1) = 6.76, p = 0.009; see Figure 1, blue bars on the left).

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

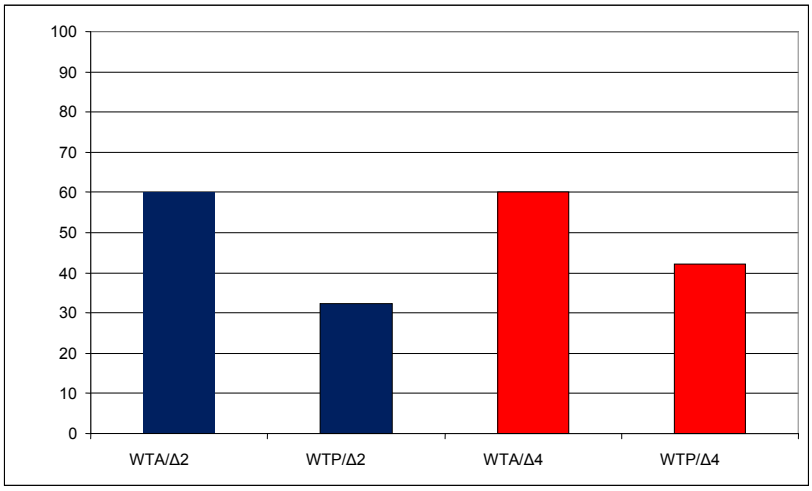

Figure 1 – Percentage of subjects who kept or switched to $10 untrackable card in Experiment 1 (vertical axis). Blue bars: Δ2 Conditions. Red bars: Δ4 Conditions.

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/Δ4] and [WTP/Δ4]): 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 not strongly significant (χ^2 (1) = 2.95, p = 0.086; Figure 1, red bars on the right). [9]

To control for age and for gender effects, we ran logistic regressions on the binary choice variable using a probit model.[10] 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 1). We used one dummy variable to control for the conditions which contrast $10 and $12 cards (Δ2=1) versus $10 and $14 cards (Δ2=0), and another dummy to control for the conditions in which the subjects were endowed with the untrackable card and were offered to

[9] Ten subjects who gave contradictory answers to the follow-up valuations questions were conservatively excluded from the analysis. Our results, however, are quite robust also to the inclusion of those subjects in the analysis set. The difference between the Δ2 conditions remains significant at the 5% level (χ^2 (1) = 4.36; p = 0.037); the difference between the Δ4 conditions remains in the hypothesized direction, but becomes less significant (χ^2 (1) = 2.52, p = 0.11). The significance levels in the regression presented further below remain the same.

[10] Differences in privacy sensitivities due to gender (Sheehan [1999]) and age (Culnan [1995]) have been observed in the literature, with younger individuals and males somewhat less prone to exhibit concern about privacy.

accept more money to switch to the tracked card (WTA=1). Age is a discrete variable and gender is a binary dummy (1=female).

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

Constant	0.9853***
	(0.32225)
Age	-0.0185***
	(.0065)
Gender	-0.0235
	(0.1962)
WTA	-0.6093***
	(0.1942)
Δ2	0.1105
	(0.1954)
	N = 179
	Prob > chi2(4) = 0.0008
	Pseudo R^2 = 0.0766

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

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. Importantly, there is no difference between cards worth $12 and those worth $14.

The initial endowment represented a form of framing**:** whether the subject wants to "sell" her data for $2 or $4 going from an untrackable to a trackable card (WTA conditions), or whether the subject wants to "pay" $2 or $4 to protect his data by switching from the trackable to the untrackable card (WTP conditions). If privacy preferences were stable and consistent, the percentages of people choosing the trackable versus the untrackable card should be the same. Instead, when subjects start with an untrackable card, they are less likely to end up with a trackable card than if they start with a trackable card.[11] Privacy valuations are therefore heavily dependent upon card endowment – a result that supports Hypothesis 1. Notably, there was no effect of the difference in card values (i.e. Δ$2 vs. Δ$4) on subjects' card choice. This result is

[11] 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, as shown in Section 2.3.

somewhat surprising, in that it points to an almost binary attitude towards privacy that is powerfully affected by WTA and WTP, but not by monetary differences. We discuss this issue further in Section 3.3.

By showing that privacy decisions respond to a non-normative factor (WTA/WTP) but not to a normative one (the objective cost associated with obtaining greater privacy), these results raise questions about the "true" value that individuals assign to the privacy of their data. Subjects who start from a situation with greater privacy protection seem to be willing to forego money to preserve their privacy. Less so, those who start from a situation of lower protection. However, since within *each* condition roughly the same proportion of subjects (60%) preferred to keep the card with which they had been endowed (regardless of its value and its privacy features), one might also wonder whether our results could simply be explained on the basis of either material or psychological transactions costs – the costs associated with departing from a default. Transactions cost are, in turn, one possible explanation for the *status quo* bias (Samuelson and Zeckhauser [1998]), the other being loss aversion. According to the transactions cost account, switching from a default can be costly (in terms of time, money, or increased uncertainty), which could make the *status quo* a normatively defensible choice, despite its inferiority to alternative options. Since Experiment 1 was based on a hypothetical choice, however, the cost incurred by 'switching' cards was negligible, so the transaction cost account cannot explain the result. Moreover, in Experiment 2, we used real gift cards, which would imply higher transaction costs. Yet, as we show below, Experiment 2 produced greater switching and, more importantly, significant differences in the proportion of subjects who switched cards depending on the experimental conditions.

3.2 Experiment 2

Whereas Experiment 1 involved hypothetical choices, Experiment 2 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. Since there was no effect of the difference in value between the trackable and untrackable cards in Experiment 1 (i.e. Δ \$2 vs. Δ \$4), in Experiment 2 we only used \$12 (trackable) and \$10 (untrackable) cards.

Subjects were shoppers at a Pittsburgh shopping mall who were offered gift cards in exchange for participating in a survey. The survey was a decoy, simply intended to create a credible reason for giving the subjects a reward (the gift card), and was identical across all conditions. Similar to

Experiment 1, 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 and two "choice" conditions. 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. 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." The choice conditions were included to situate the impact of the WTA and WTP conditions relative to a more neutral condition that did not incorporate a *status quo*. Furthermore, we included two choice conditions, one in which the anonymous $10 card appeared first, and the other in which the identified $12 card appeared first, to test for order effects. We also had 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)

3.2.1 Procedure

Experiment 2 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. Our results (and in particular the card selection) were *not* affected by the time of day when the experiment was ran, the store in front of which subjects were recruited, or whether the unrelated survey was completed before or after entering the store.

There were five different coupons, each corresponding to a study condition (see Appendix B). 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.

After completing the survey and upon exiting the store, the 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. This was true of all subjects regardless of their experimental condition.

Because the $10 and $12 gift cards looked identical, 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 this moment in the

experiment, 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.

Next, subjects in the endowed conditions were given a sheet that only contained the description of the features of the card with which they were to be endowed. The experimenter then directed the subject to select her card from the appropriate bin, be it the $10 or $12 gift card bin. In the $12 endowed, identified condition, the experimenter wrote down 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 wrote down the card number and security code of the new, $12 identified card. Therefore, the endowment manipulation was very brief, and hence, conservative, beginning after the subject chose a card from the bin, and lasting only the few seconds it took to describe the features of the card, before the subject was then asked whether she would like to exchange her card. This implies that, across conditions, subjects had a similar amount of time to reflect on how to use their respective cards *in the future* (for instance, regardless of their experimental condition, they could have compared choosing the trackable card in order to purchase non-sensitive items, versus choosing the anonymous card in order to purchase more privacy-sensitive items).

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 being manipulated between-subjects. Subjects then indicated which card they would like, and selected their card from the appropriate bin. The experimenter wrote down 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.

3.2.2 Results

Three-hundred and forty-nine female subjects participated in the study. Their mean age was 35 years and the median was 30 years. The average and median income level was $40,001-$50,000 per year; however, the modal response (16.4% of subjects) was $0-$10,000. The sample was predominantly Caucasian (83.6%). The second most common ethnicity was African American (8.5%). There were no significant differences in demographics 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%) (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 were (χ^2 (4) = 2.35, p = 0.673).

Gift card choice. Card choice did not differ as a function of whether people completed the spending survey upon entry or exit of the store (χ^2 (1) = 0.004, p = 0.948). Virtually everyone in the clarification 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 (χ^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 2).

In other words, 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%) decided to pay $2 dollars for privacy (WTP), by switching from the $12 identified card to the $10 anonymous card, to protect the same data. These results support Hypothesis 1.

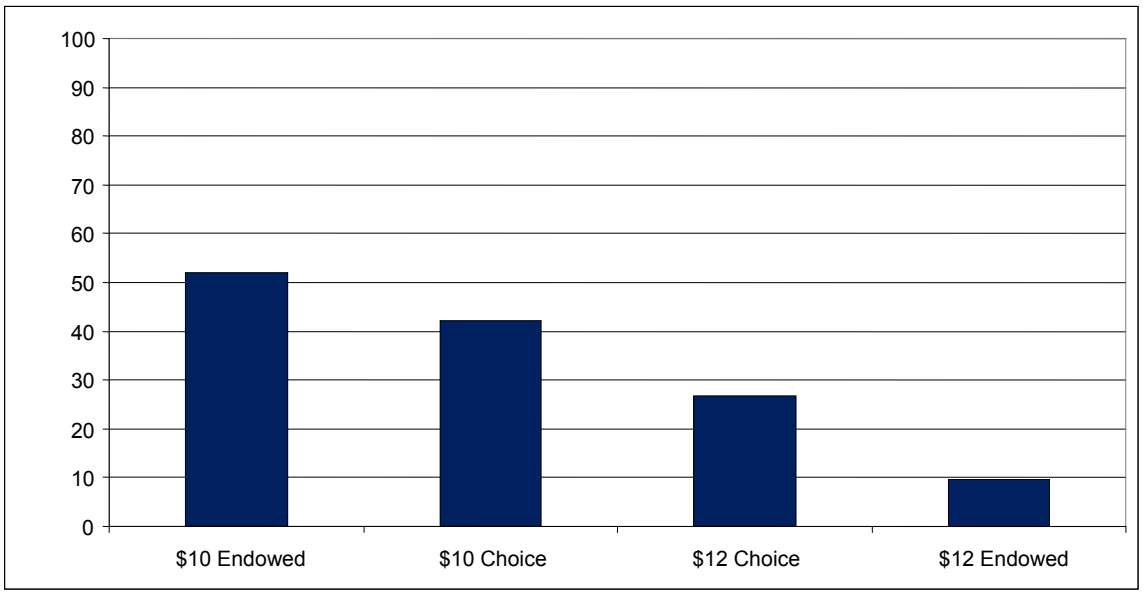

Figure 2 - Percentage of subjects who chose, kept, or switched to the $10 anonymous card in Experiment 2 (vertical axis).

These findings also help to rule out an alternative explanation for the result found in Experiment 1 (namely, that the difference in WTP and WTA may be due *status quo* bias, with the subjects in the two endowed conditions simply sticking to their default card assignment). In the real gift card experiment, subjects in the endowed conditions displayed a tendency to stick with the card to which they had been endowed (as in Experiment 1); however, while 90.3% of subjects in the $12 endowed condition chose to keep the $12 card, only 52.1% of those in the $10 endowed condition chose to keep the $10 card; in other words, significantly more subjects in the $12 endowed condition chose to keep their card than those in the $10 endowed condition χ^2 (1) = 27.24, p < 0.0005).

Ruling out a preference for default settings as an explanation for the results is also supported by contrasting the endowment conditions to the choice conditions. The two choice conditions – in which only the listed order of the card descriptions varies – are (marginally) significantly different from each other (χ^2 (1) = 3.64, p = 0.056): people are 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. However, and more relevant to our discussion, when we compare the choice conditions to the endowment conditions, we observe that subjects are significantly more likely to keep an anonymous card (instead of switching to a trackable card: 52.1% did so in the $10

endowed condition) than to *choose* an anonymous card (42.2% and 26.7% did so in the $10 and $12 choice conditions respectively: three-way Pearson χ^2 (2) = 8.76, p = 0.013). Similarly, subjects are significantly less likely to swap from a trackable to an anonymous card (9.7% did so in the $12 endowed condition) than to *choose* an anonymous card ($10 and $12 choice conditions: three-way Pearson χ^2 (2) = 18.72, p < 0.0005).

Table 2 - Probit regression, Experiment 2. The dependent variable represents the card selection (0=$10 anonymous card, 1= $12 identified card)

Constant	1.7259***
	(0.3822)
Age	-0.0178***
	(0.0063)
WTA	-0.9698***
	(0.1812)
EndORChoice	-0.0555
	(0.1724)
	N = 251
	Prob > chi2(3) = 0.0000
	Pseudo R^2 = 0.12

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

As further evidence in support of our hypotheses, we ran a logistic regression on the binary choice variable using a probit model (see Table 2). We combined the four conditions and regressed age and dummy variables representing the conditions over a dichotomous dependent variable, representing the selection of the traditional $12 gift card (1) over the privacy enhanced $10 gift card (0). We used one dummy (EndORChoice) to control for the choice (1) or endowed (0) conditions, and another dummy (WTA) to control for which card the subject was presented or endowed with first – the $10 card (1) or the $12 card (0). The model is significant, and WTA is strongly significant and negative.

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 χ^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 χ^2 (1) = 0.16, p

= 0.688), or whether the subject had been initially assigned an anonymous or identified card (Pearson χ^2 (1) = 1.28, p = 0.258), did not have an impact on their likelihood of using the card.

As an exploratory analysis, we tried to ascertain whether subjects used their cards at different types of stores, depending on card identifiability. We did not have a strong prediction, however: 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. In fact, the latter hypothesis received some support. We classified purchases depending on the store information as potentially privacy sensitive (lingerie stores such as "Victoria's Secret") or not (all other cases, including cafes such as "Aladdin's Eatery" and drugstores). We found some suggestive, albeit 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.

Discussion. Similar to the results of Experiment 1, subjects in Experiment 2 chose different gift cards depending on the framing of the choice, and therefore implicitly assigned dramatically different values to the privacy of their data. 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 in that context, 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 $2. By contrast, fewer than 10% of subjects in the identified $12 endowed condition gave up $2 to protect future purchase data. In other words, 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. Such findings imply that the *mean* valuation of the privacy of one's future purchase data differed significantly across conditions, even though subjects were randomly assigned to them. Furthermore, 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 the notion that there is a single *true* valuation of privacy to be captured.

One additional decision making factor to consider is the cost or value of a private card *relative* to the monetary amount with which subjects were initially endowed. 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 to the consideration of relative estimations of the value of privacy.

Finally, we note the remarkable difference between claimed behaviors in the hypothetical study (Experiment 1) and actual choices in the field study (Experiment 2). Roughly 40% of subjects in Experiment 1 claimed that they would switch from a $12 to a $10 card to protect their privacy, but only around 10% actually did so in Experiment 2. While we refrain from making conclusive statements about such differences (given the different methodology of the two experiments), they provides suggestive evidence that individuals may overstate their commitment to privacy choices in hypothetical conditions (see also List and Shogren [1998] on the difference between hypothetical and real choices in experiments).

3.3 The distribution of privacy valuations

Experiments 1 and 2 suggested that that the initial endowment (and the associated framing) of a privacy question as "protecting" or "revealing" personal information can have a significant impact on subjects' valuations of their data. However, a surprising finding from Experiment 1 was that there seemed to be 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 χ^2 (1) = 0.00, p = 1). Similarly, we found 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 43.2% in the latter case claimed they would switch; Pearson χ^2 (1) = 0.91, p = 0.339).

One explanation for these findings is that the valuation of the protection of purchase data in the context of the experiment simply does not vary much in the $2/$4 interval: 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 1. According to this interpretation of our findings, we could then conjecture that privacy valuations are not uniformly or even normally distributed, but are in fact clustered around some focal, extreme points. Experiment 1 provided us with additional data to test such conjecture.

In that experiment, subjects were asked – on the first page of their questionnaire – to choose between a trackable or untrackable card. These one-shot selections, alone, did not provide sufficient elements to identify exact valuation points for the subjects' privacy.[12] However, after completing their selection on the first page of the questionnaire, subjects were instructed to turn the questionnaire's page and answer a number of follow-up questions. The questions related to the valuation of the card the subject initially chose to keep or they exchange - for instance: "Would you have also kept the card you were initially given if it were an $8 card?" Such follow-up questions were designed to determine more precise valuations for data privacy; we used answers to those questions to ascertain individuals' point-wise valuations of private data.

The follow-up questions depended on the subject's answer to the one-shot question on the first page, and incremented (or decremented) by as little as 25c or as much as a few dollars. 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?"[13]

Based on the answers the subjects provided to the follow-up questions, we constructed a variable representing "brackets" of privacy valuations – namely, the approximate monetary range that

[12] For instance, a subject that keeps an untrackable $10 card rather than switching to a trackable $12 card values her privacy at least $2 – but possibly much more, or perhaps just a little bit more. Similarly, a subject who exchanges her $10 untrackable card for a trackable $12 card values her card data privacy less than $2 – but the actual value could be as little as $0, or perhaps as much as $1.99.

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

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 already know that they must value the privacy of their transaction data at least $2. Among them, now consider the person who went on to indicate 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).[14] We then took the lower boundary of each bracket, and constructed the histograms presented in Figure 4 (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).[15]

Figure 3 suggests a U shaped distribution of privacy valuations: subjects' valuations cluster at the extreme values ("between 0 and 25 cents" or "larger than $11"), with more evenly distributed valuations in between across those values, and a (local) peak at $2.[16] This U-shape can be recognized across all conditions, although it is more accentuated in the conditions where subjects were endowed with the privacy enhanced card (in the conditions in which subjects were first endowed with the more valuable and trackable card we also detected the highest concentration of low valuations for the privacy enhanced card). First, we used non-parametric ranksum 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$, $p < 0.0005$) as well as the two $10 vs. $14 conditions ($z = 2.647$, $p = 0.008$). In both cases, the conditions endowed with the more valuable but not protected card tend to assign less value to the privacy enhanced card, with is consistent with the WTP/WTA results presented in the previous sections. Then, we used Shapiro-Wilk, Shapiro-Francia, and Skewness-Kurtosis tests on the bracket data. All these tests strongly rejected the hypothesis of normality of distribution of valuations ($p <$

[14] 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.

[15] Eighteen participants did not complete the follow-up questions, and nine subjects gave irrational or inconsistent answers (i.e., accepting dominated offers and rejecting dominant offers). They were excluded from the analysis. A separate set of additional follow-up questions was not used in the analysis, since it was not consistently completed by participants.

[16] Naturally, this distribution of values is in part a function of the response options subjects were presented with, which were not evenly spaced.

0.0005).[17] Finally, Hartigan and Hartigan (1985)'s dip test for unimodality also rejected the hypothesis that the distribution is unimodal ($p < 0.0005$).[18] These results suggest that individuals' privacy valuations, in addition to being susceptible to framing, are also driven, in part, towards extreme values – possibly by idiosyncratic, subjective preferences. This distribution of valuations would be consistent with the results of Westin (1991)'s survey on privacy concerns, which 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).

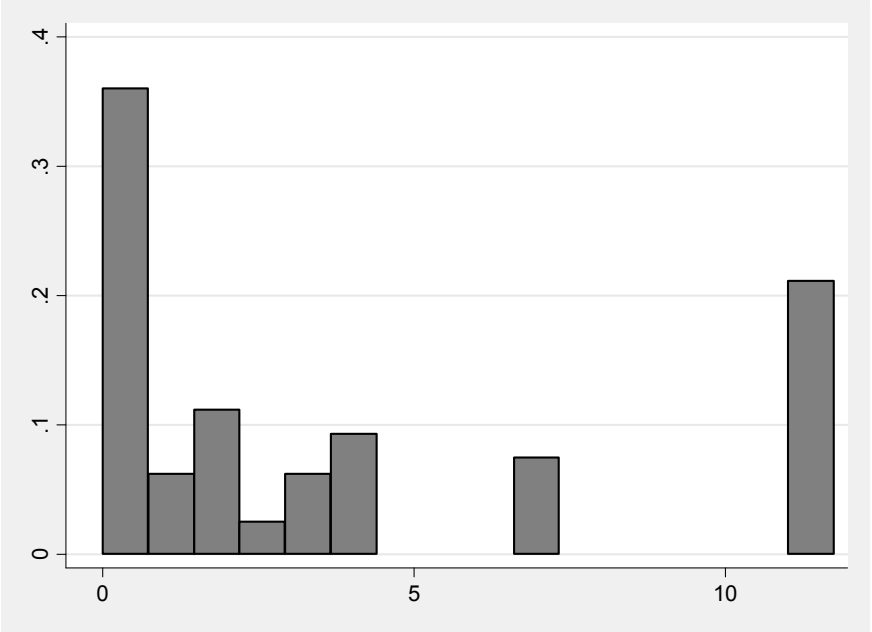

Figure 3 – Distribution of point-wise valuations of purchase data protection based on the results of Experiment 1. The vertical axis represents the fraction of observations in each range of valuations (hence, the sum of heights equals 1). The horizontal axis represents the lower boundary (in dollar terms) of each valuation bracket, from $0 to $11.

[17] The p values refer to the test applied to entire dataset. However, the hypothesis of normality is also rejected at $p < 0.05$ when considering the four conditions separately.
[18] The null hypothesis of unimodality was strongly rejected for Conditions [$10 Endowed] and [$10 Choice] ($p < 0.0005$), but was not rejected for Conditions [$12 Endowed] and [$10 Choice] ($p = 0.26$ and $p = 0.11$ respectively). The results we present in the text refer to the aggregated conditions.

4. IMPLICATIONS

Our experiments show that the minimum price a consumer will be willing to accept to allow her data to be revealed is higher than the maximum price she will be willing to pay to avoid having her data revealed – as if consumers felt "endowed" with the protection of their information, even when it refers to *future* purchase data. Our findings therefore suggest that privacy valuations, while not completely arbitrary, are subject to subtle framing effects. Specifically, 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.

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, possibly 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, our subjects may have had different motivations for opting for one card versus the other, and therefore different valuations of the protection of their data. However, thanks to randomization, subjects with different motivations – and valuations – would be similarly distributed across experimental conditions.[19] In their paper on coherent arbitrariness, Ariely, Loewenstein, and Prelec (2003) noted that their results implied 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' evaluation of personal privacy purely from market data: what people say their data is worth depends critically on the context in which they are asked - specifically, how the problem is framed.

Showing that privacy valuations may be malleable to non-normative factors is important for several reasons. First, the research is of theoretical interest because it points to the need to distinguish between decisions to protect and decisions to reveal data (whereas the IS and economics literature on privacy, so far, has assumed comparability in the behavior of individuals with respect to both types of decisions). Second, the research raises doubts about individuals'

[19] 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.

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 – such as "asymmetric" or soft forms of paternalism (Loewenstein and Haisley [2007]). Third, the finding has policy implications – similar to the ramifications of WTP/WTA dichotomies highlighted by Knetsch (1990) in the context of environmental policies. 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 valuations in the context of the design of privacy policies. The observation that individuals give away their personal information for small rewards, for example, has permeated the policy debate and has been used to argue against privacy regulation (e.g., Rubin & 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. If individual decisions regarding privacy are malleable to non-normative factors, then such arguments lose their normative standing. In our experiments, subjects who started from positions of greater privacy protection were more likely to forego money to preserve that protection.

In short, these findings suggest that we need to pay close attention to how privacy valuations are measured, whenever we want to inform the debate on privacy legislation and privacy self-regulation. "What is privacy worth?" and "Do people really care for privacy?" are questions whose answers depend not just on whom, but *how*, you ask.

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 R. Gross 2006. "Imagined Communities: Awareness, Information Sharing, and Privacy on the Facebook," *Proceedings of Privacy Enhancing Technologies Workshop (PET '06)*, 2006.

Acquisti, A. and J. Grossklags, 2005. "Privacy and rationality in decision making," *IEEE Security & Privacy*, January-February: 24-30.

Acquisti, A. and H. Varian, 2005. "Conditioning Prices on Purchase History," *Marketing Science*, 24(3), 1-15.

Ariely, D., G. Loewenstein, and D. Prelec, 2003. "Coherent Arbitrariness: Stable Demand Curves Without Stable Preferences," *Quarterly Journal of Economics*, 118(1), 73-105.

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 (1983), 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.

Culnan, M.J. and R.J. Bies, 2003. "Consumer Privacy: Balancing Economic And Justice Considerations," *Journal of Social Issues,* 59(2), 323-342.

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.

Eisenberger, R. and M. Weber,1995. "Willingness-To-Pay And Willingness-To-Accept For Risky And Ambiguous Lotteries," *Journal of Risk and Uncertainty*, 10(3), 223-233.

Fishbein, M., and I. Ajzen, 1975. *Belief, Attitude, Intention, and Behavior: An Introduction to Theory and Research*. Reading, MA: Addison-Wesley.

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." Techical 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.

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.

Knetsch, J.L., 1989. "Environmental Policy Implications Of Disparities Between Willingness To Pay And Compensation Demanded Measures Of Values," *Journal of Environmental*

Economics and Management, 18(3), 227-237.

Knetsch, J.L., 1989. "The Endowment Effect and Evidence of Nonreversible Indifference Curves," *American Economic Review*, 79(5), 1277-1284.

Knutson, B., G. Wimmer, S. Rick, N. Hollon, D. Prelec, and G. Loewenstein, 2008. "Neural Antecedents of the Endowment Effect," *Neuron*, 58(5), 814-822.

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.

Loewenstein, G. and E.C. Haisley, 2007. "The Economist as Therapist: Methodological Ramifications of 'Light' Paternalism," available at SSRN: http://ssrn.com/abstract=962472

Noam, E.M., 1996. "Privacy and self-regulation: Markets for electronic privacy," in *Privacy and Self-Regulation in the Information Ag*e, 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), 100-26.

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.

Posner, R. A., 1981. "The economics of privacy," in *American Economic Review*, 71, 405-409.

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).*

Roth, G., 2005. "Predicting the Gap between Willingness to Accept and Willingness to Pay,"

Doctoral Dissertation, Ludwig-Maximilians-Universität München, Volkswirtschaftliche Fakultät.

Rubin, P.H. and T. M. Lenard, 2002. *Privacy and the Commercial Use of Personal Information.* The Progress & Freedom Foundation, Washington, DC, USA.

Samuelson, W. and R. Zeckhauser, 1988. "Status Quo Bias In Decision Making," *Journal of Risk and Uncertainty*, 1, 7-59.

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.

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.

Syverson, P., 2003. "The Paradoxical Value Of Privacy," *Proceedings of the Second Annual Workshop on Economics and Information Security (WEIS '03),* College Park, MD.

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., 2004a. "Consumer Privacy And The Market For Customer Information," *RAND Journal of Economics*, 35(4), 631-650.

Taylor, C.R., 2004b. "Privacy And Information Acquisition In Competitive Markets*," Technical report*, *Duke University, Economics Department*.

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.

Warren S.D. and L.D. Brandeis, 1890. "The Right to Privacy," Harvard Law Review, 4(5).

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.