Losses, Gains, and Hyperbolic Discounting: An Experimental Approach to Information Security Attitudes and Behavior

Alessandro Acquisti
acquisti@sims.berkeley.edu

Jens Grossklags
jensg@sims.berkeley.edu

May 2003

UC Berkeley
2nd Annual Workshop on "Economics and Information Security"

Abstract

Surveys and experiments have uncovered a dichotomy between stated attitudes and actual behavior of individuals facing decisions affecting their privacy and their personal information security. Surveys report that most individuals are concerned about the security of their personal information and are willing to act to protect it. Experiments reveal that very few individuals actually take any action to protect their personal information, even when doing so involves limited costs. In this paper we analyze the causes of this dichotomy. We discuss which economic considerations are likely to affect individual choice and we advance testable hypotheses about why individuals' information security attitudes seem inconsistent with their behavior. We then outline an experimental design to test our hypotheses. The experiment is designed to compare individuals' characteristics as market agents to their information security attitudes and behavior.

Keywords: Information security, Privacy, Experimental Economics, Consumer Behavior.

1 Introduction

Many surveys have identified personal information security and privacy as some of the most pressing concerns of those using new information technology. On the Internet, sales for billions of dollars are said to be lost every year

because of information security fears.[1] At the same time, several technologies have been made available to protect individuals' personal information and privacy in almost any conceivable scenario - from browsing the Internet to purchasing on- and off-line. With some notable exceptions, very few of these technologies have been successful in the marketplace. There is apparently a demand, and there is an offer. So, why does market clearing seem to be absent?

In this paper we discuss which factors play a role in the decision process of individuals with respect to their information security concerns. First, we analyze economic aspects of the market for personal information security and privacy (Section 2) and advance hypotheses about why personal information attitudes seem to differ from actual behavior (Section 3). Then, we describe an experimental design to test our hypotheses. The experiment is designed to compare individuals' characteristics as market agents to their information security attitudes and behavior, and to disentangle the factors that may cause the discrepancies between the latter. (Section 4). We conclude the paper by discussing the next phases of our research (Section 5).

Our research is relevant to the formulation of information policies and to the design of information technologies for personal information security and privacy. Personal information security technologies have produced lackluster economic results in the marketplace. This signals the need to incorporate more accurate models of users' behavior into the formulation of both policy and technology. In this paper we try to offer some insights on such models.

[1] See, for example, [14].

2 Personal Information Security and Privacy: Attitudes versus Behavior

Advancements in information technology have often created new opportunities for use and risks for misuse of personal information. Recently, digital technologies and the diffusion of the Internet have have caused both popular concerns and market-based offerings of protective technologies to grown.

Rising concerns have been documented by several surveys and over time. In a Jupiter survey conducted in Spring 1999, forty percent of the 2,403 respondents said that they would have shopped on-line more often if more security of personal information could be guaranteed. A PriceWaterhouse-Coopers study in 2000 showed that nearly two thirds of the consumers surveyed abandoned more than once an on-line purchase because of privacy concerns. A Federal Trade Commission (FTC) study reported in 2000 that sixty-seven percent of consumers were "very concerned" about the privacy of the personal information provided on-line ([14]). A February 2002 Harris Interactive Survey ([22]) stated that the three biggest consumer concerns in the area of on-line personal information security were: companies trading personal data without permission, the consequences of insecure transactions, and theft of personal data. According to a Jupiter study in 2002, "$24.5 billion in on-line sales will be lost by 2006 - up from $5.5 billion in 2001. On-line retail sales would be approximately twenty-four percent higher in 2006 if consumers' fears about privacy and security were addressed effectively." ([29]).

In addition, some of the numerous surveys in this field not only reveal that individuals are concerned about the privacy and security of their personal information. They also document that certain individuals *claim* they

would be willing to take steps to protect their own information - including, in some cases, paying for it.[2]

However, more recent surveys, anecdotal evidence, and experiments have painted a different picture. [12], [21], [30], and [29] have found evidence that even privacy concerned individuals are willing to trade-off privacy for convenience or to bargain the release of very personal information in exchange of relatively small rewards. In addition, the failure of several online services aimed to provide anonymizing services to Internet users[3] provides indirect anecdotal evidence of the reluctance of most individuals to pay to protect their personal information.

Comparing these apparently conflicting data triggers three related questions:

1. Are the two sets of evidence (attitudes revealed in surveys and behavior exposed in experiments) *truly* in contradiction? In other words, is there an actual dichotomy between attitudes and behavior with regard to privacy and security of personal information - or, rather, those apparent discrepancies can be attributed to wrongful measurements and procedures?

2. If a dichotomy actually exists, what are its causes? For example, can we find a relationship between how informed an individual is about personal information security issues and her attitudes and behavior in this area? What are the relations between her market behavior as an economic agent and her behavior in terms of information security? What are the factors that ultimately determine the behavior of

[2]See Truste-Boston Consulting Group 1997 privacy survey, quoted by the Center for Democracy and Technology, www.cdt.org. Also, see www.pguardian.com, unpublished internal surveys.

[3]See [10].

information security concerned individuals?

3. Are individuals, in contrast, acting in their best interest when they choose *not* to protect themselves against possible information intrusions and accept to give away personal data in exchange for small rewards?

In the rest of this paper we comment on questions 1) and 3), but we focus on question 2). In particular, we describe our hypotheses about the heuristics applied by individuals facing information security-related decisions, and our strategy to test those hypotheses.

3 Exploring the Dichotomy

Is there a dichotomy, anyway?

The first question to address is whether, in fact, we should be at all surprised by the comparison of results from surveys such as the one reported by the FTC in 2000 and from experiments such as the one conducted by Spiekermann, Grossklags, and Berendt.

The apparent dichotomy could simply be explained by observing that different people act in different ways, and those who claim that their privacy is important are not those who fail to take actions to protect themselves.

However, that this unlikely is the case should be evident from the magnitudes of the results reported by both experimental and survey data. Although in different setups, the vast majority of subjects (both interviewed and tested) expressed privacy concerns *and* still traded-off privacy for other advantages (rewards, convenience, etc.). In addition, in their experiment Spiekermann, Grossklags, and Berendt controlled for individual behavior and attitudes for each experiment participant. They found that also those

individuals classified as privacy advocates would in fact reveal personal information in exchange of small rewards (see [30]).

Another argument that refutes the existence of a dichotomy relies on the difference between the two following concepts: 1) protecting one's privacy and information security, and 2) offering personal information in exchange of some reward. This argument emphasizes that the markets for protecting and for trading personal information may be related, but not interchangeable.

The observation that these two markets should not be confused is correct. However, the argument based on it discounts the evidence that many privacy-concerned individuals explicitly *claimed*, in surveys, to be willing to pay to protect their privacy - but then acted otherwise. In such case a dichotomy appears *within* the market for information protection. Furthermore, if the two markets for information protection and information trading are distinct (as well as the decision processes of the individuals in each market), then the above argument does not explain where the differences lie and what are their causes. Both protecting and revealing personal information imply monetary and immaterial costs and benefits (see Section 3.1). Our goal in this paper is precisely to explore the heuristics through which individuals weight these costs and benefits. In other words, the observation that the market and attitudes for information hiding may be different from the market and behavior for information sharing does not explain the existence of the dichotomy we discussed - it raises new aspects we must consider *in order to* explain it.

An additional argument against the existence of a dichotomy is that many individuals may in fact be endorsing a defensive strategy by *not completing* at all certain transactions.

Again, many individuals have certainly adopted this strategy to address

their privacy concern. Simply observing this, however, does not explain why such approach is also adopted in presence of protective technologies available at limited monetary or immaterial costs in the market. Our analysis instead is geared to understand why individuals decide to take different actions: completing a certain transaction without protecting their information, completing the transaction under the umbrella of some technology or policy that protects their information, or not completing the transaction at all. Why privacy concerned individuals can and do react in so many different ways is precisely what we attempt to understand by addressing question 2).

In doing so, we will touch also upon the related question 3): which individual behavior is optimal when her personal information security and privacy are at stake? However, we will only comment briefly on this point. We refer the reader to other (current, e.g., [1], and forthcoming) research for more in depth analysis of the existence and efficiency of an equilibrium in the market for personal information.

3.1 Privacy Advocates, Market Behavior, and Rational Choices

Individuals who claim to be concerned about their personal information act very differently when an information-sensitive situation actually arises. Some complete transactions without protecting personal information. Some give away information for small rewards. Some falsify the information they provide to other parties.[4] Some other avoid information risks altogether by aborting ongoing transactions while ignoring protecting technologies.

Are there common underlying factors explaining this variety of forms that the attitudes/behavior dichotomy takes? In this section we will try to address this question.

[4]See the 8th annual poll of the Graphics, Visualization, and Usability Center at the Georgia Institute of Technology, www.gvu.gatech.edu.

Dichotomies between attitudes and behavior have been found in several aspects of human psychology and studied in the social psychology literature since [24] and [15].[5]

One source of these observed differences can be attributed to the research procedures. It may be, for example, that people being interviewed feel a pressure to comply to a norm or want to satisfy the researcher or interviewer by providing what they consider as *correct* answers. It is equally often argued that questionnaires incur a strong bias when questioning for attitudes and reported behavior, for example, people may report a "better self" rather than their true values and attitudes. In experiments the so-called experimenter effect is an area of concern that leads to a bias of participants when they are imposed to *surveillance* in a controlled laboratory environment.

More generally, perceptions about a certain concept may vary from the moment when the concept is theoretically considered to the moment it is actually faced. "Privacy" in theory may mean many different things in practice. In the information security and privacy scenario that we consider, it may well be that many of the parameters affecting the decision process of the individual are perceived differently at the forecasting (survey) and operative (behavior) phases, thus leading to the variety of adopted strategies quoted above.

To understand this, let us abstract the decision process of an individual facing an information security issue when completing a certain transaction

[5]It is interesting to note, however, that one can also find research results where attitudes are causing a particular behavior ([6], [16] and [17]). And there exist examples where the reverse is true; that is, behavior causes attitudes([18], [19] and [8]).

in the following way:[6]

$$u_t = \delta \left[v_E(a), p^d(a) \right] + \gamma \left[v_E(t), p^d(t) \right] - c_t^d \quad (1)$$

where the utility u of completing the "transaction" t (the transaction being any action - not necessarily a monetary operation - involving exposure of personal information) is equal to some function of the *expected* payoff $v_E(t)$ from completing [or non completing] the transaction (possible revealing personal information), times the probability of completing the transaction [or not completing] with a certain technology d, $p^d(t)\ [1 - p^d(t)]$; plus some function of the *expected* payoff $v_E(a)$ from maintaining [or non maintaining] certain information secure and/or private during that transaction, times the probability of maintaining [or not maintaining] that information secure/private when using technology d, $p^d(a)\ [1 - p^d(a)]$; minus the cost of using the technology, c_t^d. The technology d may or may not be security and privacy enhancing.

Since the payoffs in 1 can be either positive or negative, that equation embodies the duality implicit in information security and privacy issues: there are both costs and benefits gained from revealing or from protecting personal information. Revealing your identity to the online bookstore may earn you a discount, or, viceversa, it may cost you a larger invoice because of price discrimination (see also [4]). Protecting your financial information by not divulging your credit card information online may save you future costs and hassles related to identity theft, but may shape your online experience more problematic and parsimonious.

With such a generic representation we do not imply that each individual is explicitly calculating all the included parameters. Rather, Equation 1

[6] See also [3].

is an abstraction that allows us to discuss the possible issues considered by individuals facing information-sensitive decisions, as well as the possible distortions that may affect their mental processes.

A rational agent, in theory, would adopt the strategy and technology that maximize her expected payoff in Equation 1: maybe completing the transaction with a security technology, maybe completing it without protection, maybe not completing the transaction at all. For example, the agent may consider the cost of sending an email through an anonymous MIX-net system (see [11]) or through a conventional, non-anonymous channel. The magnitudes of all the other factors listed in Equation 1 will change with the adopted technology. MIX-net systems will decrease the expected losses from privacy intrusions. Non-anonymous email systems will assure comparably higher reliability and raise the expected benefit from the transaction.

However, in practice, as we move from abstract representations to actual implementations, we realize that an economic agent will actually face an intricate web of trade-offs dominated by subjective evaluations and uncertainties (see also [2]). Because of these uncertainties, individuals might be discounting the potential (and subjectively evaluated) losses from losing control of their personal information with the unlikely probability that such an outcome will take place. In other words, they may perceive both the cost and the probability of losses as small. They may compare the resulting value with the implicit or explicit costs of using anonymizing technologies, which (although they may be both monetary and immaterial) are more certain and immediate. So, many privacy concerned individuals may nevertheless decide against protecting their own personal information, some may just avoid completing the transaction, and very few may actually adopt protective technologies.

The decision process described in Equation 1 therefore does not reduce to an issue of different privacy sensitivities. Several other factors may be playing a role, and their relevance may be realized by the individual only when she is facing an actual decision. More precisely, the decision process described in Equation 1 for an individual facing information security and privacy issues may be affected by the following factors (observed through surveys, user studies, and analysis):

1. **Limited information.** The amount of information the individual has access to: Is she aware of information security risks and what is her knowledge of the existence of protective technology?

2. **Benefits and costs.** There are several benefits and costs associated to using (or not using) protective technologies. Some are monetary (adoption and usage costs and benefits) and some immaterial (learning costs, switching costs, social stigma in using anonymizing technologies, hassles related to the immaturity of the technology, etc.).

3. **Bounded rationality.** Is the individual able to calculate the various parameters relevant to her choice, or is she rather limited by bounded rationality? Is she able to quantify costs and benefits of revealing or hiding information?

4. **Psychological distortions.** Are the individual's calculations affected by psychological distortions such as self-control problems, hyperbolic discounting, underinsurance?

5. **Ideology.** Is the individual considering other ideological factors that affect her privacy behavior? For example, does the individual believe that information protection is a right that the government should protect?

6. **Market behavior.** Is market behavior (such as propensity to risk, to gains or losses, and to bargaining) affecting her choice?

7. **Attitude/Behavior dichotomy.** The residual dichotomy between attitude and behavior that may be due, as discussed above, to the artificial nature of the survey environment.

If these factors impact the decision process of the individual, they may also cause the dichotomy between abstractly stated attitudes and actual behavior. Hence we discuss them in more detail below.

Limited information. The individual may not be at all aware of information security risks during certain transactions, or may ignore the existence of protective technologies, in which case the consideration of the parameters in Equation 1 would be completely distorted.

Gathering full information on every aspect of life is impossible. As a result individuals have to decide based upon incomplete or asymmetric information. Both concepts are well known in the economic literature: asymmetric information was scholarly first analyzed by Akerlof in his famous market for lemons ([7]). Varian discusses similar concepts in the privacy scenario ([32]). Incomplete information becomes a problem for the individual when she has to commit to an action without a full assessment of the associated privacy-risks. In our scenario, the individual may be ignorant about the risks she incurs by not protecting her personal information or about ways to protect herself. People may assume that institutions and governmental organizations are providing a secure platform for their actions.

Benefits and costs. There are several benefits and costs associated to using or not using information protective technologies. In particular, only some of the costs are monetary (and they could be either adoption

costs - fixed, or usage costs - variable). Other costs may be immaterial: learning costs, switching costs, usability costs, and social stigma when using anonymizing technologies, and may only be discovered through actual usage (see, for example, the difficulties in using privacy and encrypting technologies described in [33]). A survey participant may not be considering or realizing the existence of all these possible benefits and costs when answering abstract questionnaires.

One example of these costs is stigma. Goffman [20] defined stigma as an "attribute that is deeply discrediting" that reduces the bearer "from a whole and usual person to a tainted, discounted one." Consider, for example, the uneasiness of using stronger anonymizing or privacy enhancing technology, like encryption or onion-routing networks, which arises from the fear of judgement of others of what information or practices should be hidden from them. For example, personalized anonymization may be regarded as suspicious by governmental as well as by more community-based organizations.

Bounded rationality. Bounded rationality refers to both the inability to calculate probabilities and amounts for risks and related costs for the various possible individual strategies, but also to the inability to process all the uncertain and stochastic information related to information security costs and benefits.

Classic economic literature assumes humans to be rational in all aspects of life. However, even in situations with full information humans are not always capable of processing all data and deriving correct conclusions. As one of the first Herbert Simon incorporated constraints on the information-processing capacities of the individuals or entities (see [9]). Economic theories of bounded rationality can be constructed by modifying classical or

perfect rationality assumptions in various ways: (i) by introducing risk and uncertainty into demand and/or cost functions, (ii) by assuming that the entity has only incomplete information about alternatives, or (iii) by assuming complexity in the cost function or other environmental constraints so great as to prevent the actor from calculating the best course of action. The relation to the privacy notion discussed here is obvious. Individuals would collapse under the task of calculating their best strategies to minimize privacy risks for all possible interactions.

In the scenario we consider, when an individual is providing personal information to other parties, she loses control of her personal information. That loss of control multiplies, propagates, and persists for an unpredictable span of time. Hence, the individual is in a position of information asymmetry with respect to the party with whom she is transacting, and the value of the factors to be considered are very difficult to calculate correctly. In other words, the negative utility coming from future potential misuses of somebody's personal information is a random shock whose probability and scope are extremely variable, and the individual is likely in a condition of bounded rationality. For example, a small and apparently innocuous piece of information might become a crucial asset in the right context. In this case, the evaluation of the parameters in Equation 1 would be clearly distorted. Furthermore, an individual who is facing potential privacy intrusions is actually facing risks whose amounts are distributed between zero and possibly large (but mostly uncertain) amounts according to mostly unknown functions. Hence, the individual may not be able to quantify or calculate risks and benefits. In other words, individuals might decide not to protect themselves because the material and immaterial costs of protection, given the current technologies, are actually higher than the expected losses from pri-

vacy intrusions. Thus, the decision not to protect oneself paradoxically may be considered as a rational way to react to these uncertainties: the "discrepancies" between privacy attitudes and privacy behavior may reflect what could at most be called a "rational ignorance."[7]

Psychological distortions. Individuals have a tendency to "discount hyperbolically" or to show other behavioral distortions discussed in the economic literature (see, for example, [28]). Again, in this case, the evaluation of the parameters in Equation 1 would be distorted.

For example, individuals might impose constraints on their future behavior even if these constraints limit them in achieving maximum utility. This concept is incorporated into the literature as the self-control problem (sometimes also titled as changing tastes). McIntosh ([26]) tried to approach this puzzling problem in the following way: "The idea of self-control is paradoxical unless it is assumed that the psyche contains more than one energy system, and that these energy systems have some degree of independence from each other."

According to this idea, some economists now model individuals as multi-sided personalities, e.g. one personality as a farsighted planner and another one as a myopic doer ([31]).

The protection against one's future lack of own willpower could be a crucial aspect in providing a link between information security attitudes and actual behavior. People do want to protect themselves before information losses, but similarly to the attempt to stop smoking or the realization of planned consumption behavior, they might fail. One of the experiments reported in an earlier section of this paper already provided evidence for missing self-control (see, for details, [30]).

[7] See, in a different context, [25].

Furthermore, evidence of psychological experiments and observations suggest that human discounting is dynamically inconsistent. Ainslie found that discount functions are approximately hyperbolic ([5]). Hyperbolic discount functions are characterized by a relatively high discount rate over short horizons and a relatively low discount rate over long horizons. This discount structure sets up a conflict between today's preferences, and the preferences that will be held in the future ([23]). One can also relax from the assumption of a concrete functional form that is hyperbolic. However, it is generally agreed that intertemporal preferences take on the following form of time inconsistency: a person's relative preference for well-being at an earlier date over a later date gets stronger as the earlier date gets closer (present-biased preferences) ([27]).

Thus, individuals tend to under-discount long-term risks and losses while acting in privacy-sensitive situations. Note again the anecdotal finding of Jupiters' survey ([29]) that: "82 per-cent of online consumers are willing to provide various forms of information to shopping Websites from which they have yet to make purchases in exchange for something as modest as a 100 USD sweepstakes entry."

This is an interesting phenomenon, which can lead to consumer's exploitation by marketers who can design shopping sites benefitting from the immediate gratification and discounting failures of humans.

A related concept is underinsurance, the situation where an individual or entity has not arranged adequate insurance cover for the financial value of the property insured. Some researchers have already addressed this topic in detail, here also behavioral aspects where discussed. For example, Coate showed that simple altruism can lead to underinsurance by assigned recipients of donations if collective action among donors is only possible before

risks are realized ([13]).

An individual's propensity to underinsure herself against future losses that might incur with low probability but may impose a high risk emerges in the scenario we analyze. Consider, for example, the case of identity theft, where individuals' lack of carefulness can lead (with small probability) to the loss of important personal information like the Social Security Number that can then be used to create a false second identity to impose substantial financial harm on the individual.

Ideology. People might have the general belief that privacy is an enforced right, which should be guaranteed and not paid for. In this case, the individual is not adopting a mental process similar to the one described in Equation 1, but is taking a different approach, based on the advocacy of personal information rights. Hence, this is another possible psychological factor that may affect the behavior of information security-concerned individuals.

Market behavior. Finally, there may be a relation between the attitudes of a individual with respect to (for example) pricing and bargaining, and her attitude and behavior with respect to information security and privacy. In other words, market behavior may also affect the decision process of individuals who face information related issues. For example, do individuals who bargain a lot also profess more interest in privacy? Are they more or less likely to conform to those attitudes with their behavior?

In particular, let us define a "market-strategic" individual as one that knows that her actions will in turn impact the actions of another party (for example, a merchant) as in a game theoretical setup. So, for example, a strategic individual might refuse a good at a certain price in order to obtain a lesser price in a second offer (see [4]). A "market-myopic" individual on the other side will not be so forward-looking and will act following short-term

interest. Similarly, a "privacy-strategic" individual is one that calculates privacy benefits and risks and acts accordingly; a "privacy-myopic" individual on the other side will be the one who, even if she professes to appreciate privacy, does not take actions to protect herself (because of rational ignorance, as defined above, or because she only considers short-term factors). Can we compare these two sets of characteristics? For example, how many individuals act or think strategically with respect to the market, but myopically with respect to their privacy? In other words, the evaluation of γ and δ in Equation 1 could vary from individual to individual in correlation with their market characteristics, as exemplified by individuals' attitudes towards losses, gains, bargaining, and strategic behavior.

4 Experiment Design

Several hypotheses can be advanced to explain individual decision processes related to personal information security issues. Only an experimental setup under controlled conditions can determine which factors play a dominant role.

While we are not able to determine whether the parameters in Equation 1 are perceived differently at the forecasting (survey) and actually operative (behavior) phases, in this ongoing phase of our research we can address related issues through the experimental setup we have designed:

- Correlate personal information attitudes and behavior to the factors discussed in Section 3.1.
- Isolate the factors that affect the decision process of individuals with respect to their information security concerns.
- Try to explain the attitudes/behavior dichotomy through those factors.

In our approach, experiment participants are tested on their attitudes towards privacy and information security, their market behavior, and their actual personal information behavior.

To do so, subjects go through four phases:

1. Pre-questionnaire.
2. Market experiments.
3. Information behavior test.
4. Exit questionnaire.

We discuss the four experiment phases in the rest of this section. .

4.1 Experiment Phases

4.1.1 Pre-questionnaire

In the first phase, participants are interviewed over questions which may or may not be related to their personal information security and privacy attitudes. The non-related questions are used to avoid priming and influencing the behavior of the subjects in the rest of the experiment. The related questions are used to understand their privacy and information security attitudes, their knowledge about these issues, and their awareness of the risks. Examples of the pre-questionnaire questions include questions about the subject's stated concerns for personal information, her awareness of legislation or regulations in that area, her knowledge of the risks involved, her ideological position with regard to information protection, her stated willingness to take future actions to protect her information.

4.1.2 Market Experiments

In the second phase of the experiment, the subjects participate in computer-mediated games (sub-experiments) used to analyze the subjects' attitudes and behavior to the concepts discussed earlier: losses, gains, bargaining (for how long, for what amount, etc.), strategic versus myopic behavior, tendency to discount hyperbolically, attitude to self-insurance, and ability to self-control. For example, participants are given an initial monetary endowment and are asked to bet part of it in order to win a certain additional sum, or to insure themselves in order to avoid a potential loss during the rest of the experiment. The participants are also asked to engage in a bargaining game with computerized agents. The games are carefully designed to expose eventual behavioral distortions of the type discussed in the psychological-economic literature and in the previous sections.

4.1.3 Information Behavior Test

After the market experiments, each participant is offered various rewards in exchange for personal data (similarly to [30]). Subjects can accept or reject rewards of different magnitude that are offered in exchange for pieces of personal information they revealed during the two previous phases of the experiment, as well as other personal information such as name, address, etc. The announced use of the gathered data can be varied as a variable. For example, after each participant has made a first decision, she may be given another similar offer, this time with more information about how the data will be used and what risks she might incur because of revealing that information. This phase of the experiment is designed to measure information sensitivity through different reward offers.

4.1.4 Exit Questionnaire

At the end of the experiment, a set of concluding questions is presented to investigate information security and privacy technologies attitudes, knowledge and experience, including, in particular, questions that were not asked at the beginning of the experiment in order not to prime the subjects. The questions will explore in particular the subjects' knowledge of information protecting technologies, their previous usage of the latter, their evaluation of the costs involved, and so on.

4.2 Analysis of the Experiment

At the end of the experiment, three sets of data will be available: data about the subjects' information security and privacy attitudes and knowledge (from phases 1 and 4); data about their market behavior (from phase 2); data about their actual personal information behavior (from phase 3). Various forms of comparison between and among the three sets will be used to study the dichotomy discussed above.

In particular, comparison of data coming from the first and the third sets will be used to evaluate when a dichotomy between attitudes and behavior occurs. By contrasting these results to those from the second set, it will be possible to address whether these dichotomies are correlated to the subjects' market behavior.

Note that the above analysis is about correlations, not causal relations. In other words, we might find correlations between risk aversion and information-protective attitude but not between the former and actual behavior. This would not necessarily imply that risk aversion causes the discrepancy.

To move from simple correlation to causal relations, we will use the com-

parison between the answers to the repeated offers in phase 3, and we will apply econometrics techniques to test whether the factors can predict in a statistically significant way the actual behavior of the customer. By comparing at the same time the data coming from all sets, it will be possible to disentangle the causes of the dichotomy between personal information attitudes and behavior. For this purpose we are planning to use as a base model of behavior the one proposed in Section 3 above (see also [2]). That representation allows us to compare the various choices an individual considers when completing a transaction. In particular, we assume that the individual wants to maximize her utility through either protecting or not protecting her privacy. Based on the data gathered in first two phases of the experiment, we will compute the expected probability that an individual will decide to protect her privacy using a logit model, and compare it with the actual decision (observed in the third phase of the experiment).

Our goal here is to isolate discrepancies due to lack of information or other factors from those due to low privacy sensitivity. For example, if we find out that a subject whose privacy behavior differs from her attitudes towards privacy while the subject herself does not display behavioral distortions or ignorance of the risks, we might conclude that she actually has a lower privacy sensitivity than what she claims in the interview (that is, in phase 1 of the experiment).

5 Conclusions

In this paper we have discussed economic aspects of the market for personal information security and privacy. Privacy seems easier to protect than to "sell," in the sense that many privacy-enhancing technologies are available but few have succeeded in the market. Using economic reasoning we

have advanced some hypotheses about why privacy attitudes apparently differ from privacy behavior: limited information, self-control problems, other behavioral distortions, bounded rationality. We have described an experimental design that would allow us to verify our hypotheses and disentangle the factors that may cause this discrepancy.

The mixed results met in the marketplace by personal information security technologies is evidence of the need to incorporate more accurate models of user's behavior into the formulation of policy and technology guidelines. We hope that our ongoing analysis can be useful to the design of information policies and information technologies.

References

[1] Alessandro Acquisti. Privacy and security of personal information: Economic incentives and technological solutions. In *1st SIMS Workshop on Economics and Information Security*, 2002. http://www.sims.berkeley.edu/ acquisti/papers/.

[2] Alessandro Acquisti. Protecting privacy with economics: Economic incentives for preventive technologies in ubiquitous computing environments. In *Workshop on Socially-informed Design of Privacy-enhancing Solutions, 4th International Conference on Ubiquitous Computing - UBICOMP '02*, 2002.

[3] Alessandro Acquisti, Roger Dingledine, and Paul Syverson. Open issues in the economics of anonymity. In *Financial Cryptography - FC '03*. Springer Verlag, LNCS, forthcoming, 2003.

[4] Alessandro Acquisti and Hal R. Varian. Conditioning prices on pur-

chase history. Technical report, University of California, Berkeley, 2002. http://www.sims.berkeley.edu/~acquisti/papers/.

[5] George W. Ainslie. Specious reward: A behavioral theory of impulsiveness and impulsive control. *Psychological Bulletin*, 82:463–496, 1975.

[6] Icek Ajzen. *Attitudes, personality, and behavior*, chapter 6. Open University Press, Milton-Keynes, England, 1988.

[7] George A. Akerlof. The market for "lemons": Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84:488–500, 1970.

[8] Elliot Aronson and Judson Mills. The effect of severity of initiation on the devaluation of forbidden behavior. *Journal of Abnormal and Social Psychology*, 59:177–181, 1959.

[9] Peter L. Berger. *Models of bounded rationality, Vol. I-III.* The MIT Press, Cambridge, MA, 1982.

[10] Benjamin D. Brunk. Understanding the privacy space. *First Monday*, 7, 2002. "http://firstmonday.org/issues/issue7_10/brunk/index.html.

[11] David Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(2):84–88, 1981.

[12] Ramnath K. Chellappa and Raymong Sin. Personalization versus privacy: An empirical examination of the online consumer's dilemma. In *2002 Informs Meeting*, 2002.

[13] Stephen Coate. Altruism, the samaritan's dilemna, and government transfer policy. *American Economic Review*, 85(1):46–57, 1995.

[14] Federal Trade Commission. Privacy online: Fair information practices in the electronic marketplace, 2000. http://www.ftc.gov/reports/privacy2000/privacy2000.pdf.

[15] S.M. Corey. Professional attitudes and actual behavior. *Journal of educational psychology*, 28(1):271 – 280, 1937.

[16] Alice H. Eagly and Shelly Chaiken. *The Psychology of Attitudes*, chapter 4. Harcourt Brace Jovanovich College Publishers, Fort Worth, TX, 1993.

[17] Russell H. Fazio. Multiple processes by which attitudes guide behavior: The mode model as an integrative framework. *Advances in experimental social psychology*, 23:75–109, 1990.

[18] Leon Festinger. *A theory of cognitive dissonance.* Row Peterson, Evanston, IL, 1957.

[19] Leon Festinger and James M. Carlsmith. Cognitive consequences of forced compliance. *Journal of Abnormal and Social Psychology*, 58:203–210, 1959.

[20] Erving Goffman. *Stigma: Notes on the Management of Spoiled Identity.* Prentice-Hall, Englewood Cliffs, NJ, 1963.

[21] Il-Horn Harn, Kai-Lung Hui, Tom S. Lee, and Ivan P. L. Png. Online information privacy: Measuring the cost-benefit trade-off. In *23rd International Conference on Information Systems*, 2002.

[22] Harris Interactive. First major post-9-11 privacy survey finds consumers demanding companies do more to protect privacy; public wants company privacy policies to be independently verified, 2002. http://www.harrisinteractive.com/news/allnewsbydate.asp?NewsID=429.

[23] David Laibson. Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics*, 62(2):443–477, 1997.

[24] Robert LaPiere. Attitudes versus actions. *Social Forces*, 13:230–237, 1934.

[25] Mark Lemley. Rational ignorance at the patent office. Technical report, Berkeley Olin Program in Law and Economics, Working Paper Series, 2000.

[26] Donald McIntosh. *The Foundations of Human Society*. The University of Chicago Press, Chicago, IL, 1969.

[27] Ted O'Donoghue and Matthew Rabin. Choice and procrastination. *Quarterly Journal of Economics*, 116(1):121–160, 2001.

[28] Matthew Rabin and Ted O'Donoghue. The economics of immediate gratification. *Journal of Behavioral Decision Making*, 13(2):233–250, 2000.

[29] Jupiter Research. Seventy percent of US consumers worry about online privacy, but few take protective action, 2002. `http://www.jmm.com/xp/jmm/press/2002/pr_060302.xml`.

[30] Sarah Spiekermann, Jens Grossklags, and Bettina Berendt. E-privacy in 2nd generation e-commerce: Privacy preferences versus actual behavior. In *3rd ACM Conference on Electronic Commerce - EC '01*, pages 38–47, 2002.

[31] Richard Thaler and Hersh M. Shefrin. An economic theory of self-control. *The Journal of Political Economy*, 89:392–406, 1981.

[32] Hal R. Varian. Economic aspects of personal privacy. In *Privacy and Self-Regulation in the Information Age*. National Telecommunications and Information Administration, 1996.

[33] Alma Whitten and J. D. Tygar. Why Johnny can't encrypt: A usability evaluation of PGP 5.0. In *8th USENIX Security Symposium*, 1999. `citeseer.nj.nec.com/whitten99why.html`.