Privacy in Electronic Commerce and the Economics of Immediate Gratification

Alessandro Acquisti
H. John Heinz III School of Public Policy and Management
Carnegie Mellon University
acquisti@andrew.cmu.edu

ABSTRACT

Dichotomies between privacy attitudes and behavior have been noted in the literature but not yet fully explained. We apply lessons from the research on behavioral economics to understand the individual decision making process with respect to privacy in electronic commerce. We show that it is unrealistic to expect individual rationality in this context. Models of self-control problems and immediate gratification offer more realistic descriptions of the decision process and are more consistent with currently available data. In particular, we show why individuals who may genuinely want to protect their privacy might not do so because of psychological distortions well documented in the behavioral literature; we show that these distortions may affect not only 'naïve' individuals but also 'sophisticated' ones; and we prove that this may occur also when individuals perceive the risks from not protecting their privacy as significant.

Categories and Subject Descriptors

J.4 [**Social and Behavioral Sciences**]: Economics; K.4.1 [**Public Policy Issues**]: Privacy

General Terms

Economics, Security, Human Factors

Keywords

Privacy, Electronic Commerce, Immediate Gratification, Hyperbolic Discounting, Self-Control Problems

1. PRIVACY AND ELECTRONIC COMMERCE

Privacy remains an important issue for electronic commerce. A PriceWaterhouseCoopers study in 2000 showed that nearly two thirds of the consumers surveyed "would shop more online if they knew retail sites would not do anything with their personal information" [15]. A Federal Trade Commission study reported in 2000 that sixty-seven percent of consumers were "very concerned" about the privacy of the personal information provided on-line [11]. More recently, a February 2002 Harris Interactive survey found 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 [19]. According to a Jupiter Research study in 2002, "$24.5 billion in on-line sales will be lost by 2006 - up from $5.5 billion in 2001. Online retail sales would be approximately twenty-four percent higher in 2006 if consumers' fears about privacy and security were addressed effectively" [21]. Although the media hype has somewhat diminished, risks and costs have not - as evidenced by the increasing volumes of electronic spam and identity theft [16].

Surveys in this field, however, as well as experiments and anecdotal evidence, have also painted a different picture. [36, 10, 18, 21] have found evidence that even privacy concerned individuals are willing to trade-off privacy for convenience, or bargain the release of very personal information in exchange for relatively small rewards. The failure of several on-line services aimed at providing anonymity for Internet users [6] offers additional indirect evidence of the reluctance by most individuals to spend any effort in protecting their personal information.

The dichotomy between privacy attitudes and behavior has been highlighted in the literature. Preliminary interpretations of this phenomenon have been provided [2, 38, 33, 40]. Still missing are: an explanation grounded in economic or psychological theories; an empirical validation of the proposed explanation; and, of course, the answer to the most recurring question: should people bother at all about privacy?

In this paper we focus on the first question: we formally analyze the individual decision making process with respect to privacy and its possible shortcomings. We focus on individual (mis)conceptions about their handling of risks they face when revealing private information. We do *not* address the issue of whether people should actually protect themselves. We will comment on that in Section 5, where we will also discuss strategies to empirically validate our theory.

We apply lessons from behavioral economics. Traditional economics postulates that people are forward-looking and bayesian updaters: they take into account how current behavior will influence their future well-being and preferences. For example, [5] study *rational* models of addiction. This approach can be compared to those who see in the decision

EC'04, May 17–20, 2004, New York, New York, USA.

not to protect one's privacy a rational choice given the (supposedly) low risks at stake. However, developments in the area of behavioral economics have highlighted various forms of psychological inconsistencies (self-control problems, hyperbolic discounting, present-biases, etc.) that clash with the fully rational view of the economic agent. In this paper we draw from these developments to reach the following conclusions:

- We show that it is unlikely that individuals can act rationally in the economic sense when facing privacy sensitive decisions.
- We show that alternative models of personal behavior and time-inconsistent preferences are compatible with the dichotomy between attitudes and behavior and can better match current data. For example, they can explain the results presented by [36] at the ACM EC '01 conference. In their experiment, self-proclaimed privacy advocates were found to be willing to reveal varying amounts of personal information in exchange for small rewards.
- In particular, we show that individuals may have a tendency to under-protect themselves against the privacy risks they perceive, and over-provide personal information even when wary of (perceived) risks involved.
- We show that the magnitude of the perceived costs of privacy under certain conditions will *not* act as deterrent against behavior the individual admits is risky.
- We show, following similar studies in the economics of immediate gratification [31], that even 'sophisticated' individuals may under certain conditions become 'privacy myopic.'

Our conclusion is that simply providing more information and awareness in a self-regulative environment is not sufficient to protect individual privacy. Improved technologies, by lowering costs of adoption and protection, certainly can help. However, more fundamental human behavioral responses must also be addressed if privacy ought to be protected.

In the next section we propose a model of rational agents facing privacy sensitive decisions. In Section 3 we show the difficulties that hinder any model of privacy decision making based on full rationality. In Section 4 we show how behavioral models based on immediate gratification bias can better explain the attitudes-behavior dichotomy and match available data. In Section 5 we summarize and discuss our conclusions.

2. A MODEL OF RATIONALITY IN PRIVACY DECISION MAKING

Some have used the dichotomy between privacy attitudes and behavior to claim that individuals are acting rationally when it comes to privacy. Under this view, individuals may accept small rewards for giving away information because they expect future damages to be even smaller (when discounted over time and with their probability of occurrence). Here we want to investigate what underlying assumptions about personal behavior would support the hypothesis of full rationality in privacy decision making.

Since [28, 37, 29] economists have been interested in privacy, but only recently formal models have started appearing [3, 7, 39, 40]. While these studies focus on market interactions between one agent and other parties, here we are interested in formalizing the decision process of the single individual. We want to see if individuals can be economically rational (forward-lookers, bayesian updaters, utility maximizers, and so on) when it comes to protect their own personal information.

The concept of privacy, once intended as the right to be left alone [41], has transformed as our society has become more information oriented. In an information society the self is expressed, defined, and affected through and by information and information technology. The boundaries between private and public become blurred. Privacy has therefore become more a *class* of multifaceted interests than a single, unambiguous concept. Hence its value may be discussed (if not ascertained) only once its context has also been specified. This most often requires the study of a network of relations between a subject, certain information (related to the subject), other parties (that may have various linkages of interest or association with that information or that subject), and the context in which such linkages take place.

To understand how a rational agent could navigate through those complex relations, in Equation 1 we abstract the decision process of an idealized rational economic agent who is facing privacy trade-offs when completing a certain transaction.

$$\max_{d} U_t = \delta\left(v_E\left(a\right), p^d\left(a\right)\right) + \gamma\left(v_E\left(t\right), p^d\left(t\right)\right) - c_t^d \quad (1)$$

In Equation 1, δ and γ are unspecified functional forms that describe weighted relations between expected payoffs from a set of events v and the associated probabilities of occurrence of those events p. More precisely, the utility U of completing a transaction t (the transaction being any action - not necessarily a monetary operation - *possibly* involving exposure of personal information) is equal to some function of the *expected* payoff $v_E\left(a\right)$ from maintaining (or not) certain information private during that transaction, and the probability of maintaining [or not maintaining] that information private when using technology d, $p^d\left(a\right)$ $\left[1 - p^d\left(a\right)\right]$; plus some function of the *expected* payoff $v_E\left(t\right)$ from completing (or non completing) the transaction (possibly revealing personal information), and the probability of completing [or not completing] that transaction with a certain technology d, $p^d\left(t\right)$ $\left[1 - p^d\left(t\right)\right]$; minus the cost of using the technology t: c_t^d.[1]

The technology d may or may not be privacy enhancing. Since the payoffs in Equation 1 can be either positive or negative, Equation 1 embodies the duality implicit in privacy issues: there are both costs and benefits gained from revealing or from protecting personal information, and the costs and benefits from completing a transaction, $v_E\left(t\right)$, might be distinct from the costs and benefits from keeping the associated *information* private, $v_E\left(a\right)$. For instance, revealing one's identity to an on-line bookstore may earn a discount. Viceversa, it may also cost a larger bill, because of price discrimination. Protecting one's financial privacy by not divulging credit card information on-line may protect against future losses and hassles related to identity theft. But it may

[1] See also [1].

make one's on-line shopping experience more cumbersome, and therefore more expensive.

The functional parameters δ and γ embody the variable weights and attitudes an individual may have towards keeping her information private (for example, her privacy sensitivity, or her belief that privacy is a right whose respect should be enforced by the government) and completing certain transactions. Note that v_E and p could refer to *sets* of payoffs and the associated probabilities of occurrence. The payoffs are themselves only *expected* because, regardless of the probability that the transaction is completed or the information remains private, they may depend on other sets of events and their associated probabilities. $v_E()$ and $p^d()$, in other words, can be read as multi-variate parameters inside which are hidden several other variables, expectations, and functions because of the complexity of the privacy network described above.

Over time, the probability of keeping certain information private, for instance, will not only depend on the chosen technology d but also on the efforts by other parties to appropriate that information. These efforts may be function, among other things, of the expected value of that information to those parties. The probability of keeping information private will also depend on the environment in which the transaction is taking place. Similarly, the expected benefit from keeping information private will also be a collection over time of probability distributions dependent on several parameters. Imagine that the probability of keeping your financial transactions private is very high when you use a bank in Bermuda: still, the expected value from keeping your financial information confidential will depend on a number of other factors.

A rational agent would, *in theory*, choose the technology d that maximizes her expected payoff in Equation 1. Maybe she would choose to complete the transaction under the protection of a privacy enhancing technology. Maybe she would complete the transaction without protection. Maybe she would not complete the transaction at all ($d = 0$). For example, the agent may consider the costs and benefits of sending an email through an anonymous MIX-net system [8] and compare those to the costs and benefits of sending that email through a conventional, non-anonymous channel. The magnitudes of the parameters in Equation 1 will change with the chosen technology. MIX-net systems may decrease the expected losses from privacy intrusions. Non-anonymous email systems may promise comparably higher reliability and (possibly) reduced costs of operations.

3. RATIONALITY AND PSYCHOLOGICAL DISTORTIONS IN PRIVACY

Equation 1 is a comprehensive (while intentionally generic) road-map for navigation across privacy trade-offs that no human agent would be *actually* able to use.

We hinted to some difficulties as we noted that several layers of complexities are hidden inside concepts such as the "expected value of maintaining certain information private," and the "probability" of succeeding doing so. More precisely, an agent will face three problems when comparing the trade-offs implicit in Equation 1: incomplete information about *all* parameters; *bounded* power to process all available information; no deviation from the rational path towards utility-maximization. Those three problems are precisely the same issues real people have to deal with on an everyday basis as they face privacy-sensitive decisions. We discuss each problem in detail.

1. Incomplete information. What information has the individual access to as she prepares to take privacy sensitive decisions? For instance, is she aware of privacy invasions and the associated risks? What is her knowledge of the existence and characteristics of protective technologies?

Economic transactions are often characterized by incomplete or asymmetric information. Different parties involved may not have the same amount of information about the transaction and may be uncertain about some important aspects of it [4]. Incomplete information will affect almost all parameters in Equation 1, and in particular the estimation of costs and benefits. Costs and benefits associated with privacy protection and privacy intrusions are both monetary and immaterial. Monetary costs may for instance include adoption costs (which are probably fixed) and usage costs (which are variable) of protective technologies - if the individual decides to protect herself. Or they may include the financial costs associated to identity theft, if the individual's information turns out not to have been adequately protected. Immaterial costs may include learning costs of a protective technology, switching costs between different applications, or social stigma when using anonymizing technologies, and many others. Likewise, the benefits from protecting (or not protecting) personal information may also be easy to quantify in monetary terms (the discount you receive for revealing personal data) or be intangible (the feeling of protection when you send encrypted emails).

It is difficult for an individual to estimate all these values. Through information technology, privacy invasions can be ubiquitous and invisible. Many of the payoffs associated with privacy protection or intrusion may be discovered or ascertained only *ex post* through actual experience. Consider, for instance, the difficulties in using privacy and encrypting technologies described in [43].

In addition, the calculations implicit in Equation 1 depend on incomplete information about the probability distribution of future events. Some of those distributions may be predicted after comparable data - for example, the probability that a certain credit card transaction will result in fraud today could be calculated using existing statistics. The probability distributions of other events may be very difficult to estimate because the environment is too dynamic - for example, the probability of being subject to identity theft 5 years in the future because of certain data you are releasing *now*. And the distributions of some other events may be almost completely subjective - for example, the probability that a new and practical form of attack on a currently secure cryptosystem will expose all of your encrypted personal communications a few years from now.

This leads to a related problem: bounded rationality.

2. Bounded rationality. Is the individual able to *calculate* all the parameters relevant to her choice? Or is she limited by bounded rationality?

In our context, bounded rationality refers to the inability to calculate and compare the magnitudes of payoffs associated with various strategies the individual may choose in privacy-sensitive situations. It also refers to the inability to process all the stochastic information related to risks and probabilities of events leading to privacy costs and benefits.

In traditional economic theory, the agent is assumed to have both rationality and unbounded 'computational' power to process information. But human agents are unable to process all information in their hands and draw accurate conclusions from it [34]. In the scenario we consider, once an individual provides personal information to other parties, she literally loses control of that information. That loss of control propagates through other parties and persists for unpredictable spans of time. Being in a position of information asymmetry with respect to the party with whom she is transacting, decisions must be based on stochastic assessments, and the magnitudes of the factors that may affect the individual become very difficult to aggregate, calculate, and compare.[2] Bounded rationality will affect the calculation of the parameters in Equation 1, and in particular δ, γ, $v_E()$, and $p_t()$. The cognitive costs involved in trying to calculate the best strategy could therefore be so high that the individual may just resort to simple heuristics.

3. Psychological distortions. Eventually, even if an individual had access to complete information and could appropriately compute it, she still may find it difficult to follow the rational strategy presented in Equation 1. A vast body of economic and psychological literature has by now confirmed the impact of several forms of psychological distortions on individual decision making. Privacy seems to be a case study encompassing many of those distortions: hyperbolic discounting, under insurance, self-control problems, immediate gratification, and others. The traditional dichotomy between attitude and behavior, observed in several aspects of human psychology and studied in the social psychology literature since [24] and [13], may also appear in the privacy space because of these distortions.

For example, individuals have a tendency to discount 'hyperbolically' future costs or benefits [31, 27]. In economics, hyperbolic discounting implies inconsistency of personal preferences over time - future events may be discounted at different discount rates than near-term events. Hyperbolic discounting may affect privacy decisions, for instance when we heavily discount the (low) probability of (high) future risks such as identity theft.[3] Related to hyperbolic discounting is the tendency to underinsure oneself against certain risks [22].

In general, individuals may put constraints on future behavior that limit their own achievement of maximum utility: people may genuinely want to protect themselves, but because of self-control bias, they will not actually take those steps, and opt for immediate gratification instead. "People tend to underappreciate the effects of changes in their states, and hence falsely project their current preferences over consumption onto their future preferences. Far more than suggesting merely that people mispredict future tastes, this projection bias posits a systematic pattern in these mispredictions which can lead to systematic errors in dynamic-choice environments" [25, p. 2].

In addition, individuals suffer from optimism bias [42], the misperception that one's risks are lower than those of other individuals under similar conditions. Optimism bias may lead us to believe that we will not be subject to privacy intrusions.

Individuals encounter difficulties when dealing with cumulative risks. [35], for instance, shows that while young smokers appreciate the long term risks of smoking, they do not fully realize the cumulative relation between the low risks of each additional cigarette and the slow building up of a serious danger. Difficulties with dealing with cumulative risks apply to privacy, because our personal information, once released, can remain available over long periods of time. And since it can be correlated to other data, the 'anonymity sets' [32, 14] in which we wish to remain hidden get smaller. As a result, the whole risk associated with revealing different pieces of personal information is *more* than the sum of the individual risks associated with each piece of data.

Also, it is easier to deal with actions and effects that are closer to us in time. Actions and effects that are in the distant future are difficult to focus on given our limited foresight perspective. As the foresight changes, so does behavior, even when preferences remain the same [20]. This phenomenon may also affects privacy decisions, since the costs of privacy protection may be immediate, but the rewards may be invisible (absence of intrusions) and spread over future periods of time.

To summarize: whenever we face privacy sensitive decisions, we hardly have all data necessary for an informed choice. But even if we had, we would be likely unable to process it. And even if we could process it, we may still end behaving against our own better judgment. In what follows, we present a model of privacy attitudes and behavior based on some of these findings, and in particular on the plight of immediate gratification.

4. PRIVACY AND THE ECONOMICS OF IMMEDIATE GRATIFICATION

The problem of immediate gratification (which is related to the concepts of time inconsistency, hyperbolic discounting, and self-control bias) is so described by O'Donoghue and Rabin [27, p. 4]: "A person's relative preference for well-being at an earlier date over a later date gets stronger as the earlier date gets closer. [...] [P]eople have self-control problems caused by a tendency to pursue immediate gratification in a way that their 'long-run selves' do not appreciate." For example, if you were given only two alternatives, on Monday you may claim you will prefer working 5 hours on Saturday to 5 hours and half on Sunday. But as Saturday comes, you will be more likely to prefer postponing work until Sunday.

This simple observation has rather important consequences in economic theory, where time-consistency of preferences is the dominant model. Consider first the traditional model of utility that agents derive from consumption: the model states that utility discounts exponentially over time:

$$U_t = \sum_{\tau=t}^{T} \delta^{\tau} u_{\tau} \tag{2}$$

In Equation 2, the cumulative utility U at time t is the discounted sum of all utilities from time t (the present) until time T (the future). δ is the discount factor, with a value

[2] The negative utility coming from future potential misuses of somebody's personal information could be a random shock whose probability and scope are extremely variable. For example, a small and apparently innocuous piece of information might become a crucial asset or a dangerous liability in the right context.

[3] A more rigorous description and application of hyperbolic discounting is provided in Section 4.

	Period 1	Period 2	Period 3	Period 4
Benefits from selling period 1	2	0	0	0
Costs from selling period 1	0	1	1	1
Benefits from selling period 2	0	2	0	0
Costs from selling period 2	0	0	1	1
Benefits from selling period 3	0	0	2	0
Costs from selling period 3	0	0	0	1

Table 1: (Fictional) expected payoffs from joining loyalty program.

between 0 and 1. A value of 0 would imply that the individual discounts so heavily that the utility from future periods is worth zero today. A value of 1 would imply that the individual is so patient she does not discount future utilities. The discount factor is used in economics to capture the fact that having (say) one dollar one year from now is valuable, but not as much as having that dollar *now*. In Equation 2, if all u_τ were constant - for instance, 10 - and δ was 0.9, then at time $t = 0$ (that is, *now*) u_0 would be worth 10, but u_1 would be worth 9.

Modifying the traditional model of utility discounting, [23] and then [31] have proposed a model which takes into account possible time-inconsistency of preferences. Consider Equation 3:

$$U_t(u_t, u_{t+1}, ..., u_T) = \delta^t u_t + \beta \sum_{\tau=t+1}^{T} \delta^\tau u_\tau \tag{3}$$

Assume that $\delta, \beta \in [0, 1]$. δ is the discount factor for intertemporal utility as in Equation 2. β is the parameter that captures an individual's tendency to gratify herself immediately (a form of time-inconsistent preferences). When β is 1, the model maps the traditional time-consistent utility model, and Equation 3 is identical to Equation 2. But when β is zero, the individual does not care for anything but today. In fact, any β smaller than 1 represents self-control bias.

The experimental literature has convincingly proved that human beings tend to have self-control problems even when they claim otherwise: we tend to avoid and postpone undesirable activities even when this will imply more effort tomorrow; and we tend to over-engage in pleasant activities even though this may cause suffering or reduced utility in the future.

This analytical framework can be applied to the study of privacy attitudes and behavior. Protecting your privacy sometimes means protecting yourself from a clear and present hassle (telemarketers, or people peeping through your window and seeing how you live - see [33]); but sometimes it represents something akin to getting an insurance against future and only uncertain risks. In surveys completed at time $t = 0$, subjects asked about their attitude towards privacy risks may mentally consider some costs of protecting themselves at a later time $t = s$ and compare those to the avoided costs of privacy intrusions in an even more distant future $t = s + n$. Their alternatives at survey time 0 are represented in Equation 4.

$$\min_{\text{wrt } x} DU_0 = \beta[(E(c_{s,p})\delta^s x) + (E(c_{s+n,i})\delta^{s+n}(1-x))] \tag{4}$$

x is a dummy variable that can take values 0 or 1. It represents the individual's choice - which costs the individual opts to face: the expected cost of protecting herself at time s, $E(c_{s,p})$ (in which case $x = 1$), or the expected costs of being subject to privacy intrusions at a later time $s + n$, $E(c_{s+n,i})$.

The individual is trying to minimize the disutility DU of these costs with respect to x. Because she discounts the two future events with the same discount factor (although at different times), for certain values of the parameters the individual may conclude that paying to protect herself is worthy. In particular, this will happen when:

$$E(c_{s,p})\delta^s < E(c_{s+n,i})\delta^{s+n} \tag{5}$$

Now, consider what happens as the moment $t = s$ comes. Now a real price should be paid in order to enjoy some form of protection (say, starting to encrypt all of your emails to protect yourself from future intrusions). Now the individual will *perceive* a different picture:

$$\min_{\text{wrt } x} DU_s = E(c_{s,p})x + \beta E(c_{n,i})\delta^n(1-x) \tag{6}$$

Note that nothing has changed in the equation (certainly not the individual's perceived risks) except *time*. If β (the parameter indicating the degree of self-control problems) is less than one, chances are that the individual now will actually choose *not* to protect herself. This will in fact happen when:

$$E(c_{s,p}) > \beta E(c_{n,i})\delta^n \tag{7}$$

Note that Disequalities 5 and 7 may be simultaneously met for certain $\beta < 1$. At survey time the individual honestly claimed she wanted to protect herself *in principle* - that is, some time in the future. But as she is asked to make an effort to protect herself right now, she chooses to run the risk of privacy intrusion.

Similar mathematical arguments can be made for the comparison between immediate costs with immediate benefits (subscribing to a 'no-call' list to stop telemarketers from harassing you at dinner), and immediate costs with only future expected rewards (insuring yourself against identity theft, or protecting yourself from frauds by never using your credit card on-line), particularly when expected future rewards (or avoided risks) are also intangible: the immaterial consequences of living (or not) in a dossier society, or the chilling effects (or lack thereof) of being under surveillance.

The reader will have noticed that we have focused on *perceived* (expected) costs $E(c)$, rather than real costs. We do not know the real costs and we do not claim that the individual does. But we are able to show that under certain conditions even costs perceived as very high (as during periods of intense privacy debate) will be ignored.

We can provide some fictional numerical examples to make the analysis more concrete. We present some scenarios inspired by the calculations in [31].

Imagine an economy with just 4 periods (Table 1). Each individual can enroll in a supermarket's loyalty program by revealing personal information. If she does so, the individual gets a discount of 2 during the period of enrollment, only to pay one unit each time thereafter because of price discrimination based on the information she revealed (we make no attempt at calibrating the realism of this obviously abstract example; the point we are focusing on is how time inconsistencies may affect individual behavior given the expected costs and benefits of certain actions).[4] Depending on which period the individual chooses for 'selling' her data, we have the undiscounted payoffs represented in Table 1.

Imagine that the individual is contemplating these options and discounting them according to Equation 3. Suppose that $\delta = 1$ for all types of individuals (this means that for simplicity we do not consider intertemporal discounting) but $\beta = 1/2$ for time-inconsistent individuals and $\beta = 1$ for everybody else. The time-consistent individual will choose to join the program at the very last period and rip off a benefit of 2-1=1. The individual with immediate gratification problems, for whom $\beta = 1/2$, will instead perceive the benefits from joining now or in period 3 as equivalent (0.5), and will join the program now, thus actually making herself worse off.

[31] also suggest that, in addition to the distinction between time-consistent individuals and individuals with time-inconsistent preferences, we should also distinguish time-inconsistent individuals who are *naïve* from those who are *sophisticated*. Naïve time-inconsistent individuals are not aware of their self-control problems - for example, they are those who always plan to start a diet *next week*. Sophisticated time-inconsistent individuals *suffer* of immediate gratification bias, but are at least *aware* of their inconsistencies. People in this category choose their behavior today correctly estimating their future time-inconsistent behavior.

Now consider how this difference affects decisions in another scenario, represented in Table 2. An individual is considering the adoption of a certain privacy enhancing technology. It will cost her some money both to protect herself and *not* to protect herself. If she decides to protect herself, the cost will be the amount she pays - for example - for some technology that shields her personal information. If she decides not to protect herself, the cost will be the expected consequences of privacy intrusions.

We assume that *both* these aggregate costs increase over time, although because of separate dynamics. As time goes by, more and more information about the individual has been revealed, and it becomes more costly to be protected against privacy intrusions. At the same time, however, intrusions become more frequent and dangerous.

In period 1, the individual may protect herself by spending 5, or she may choose to face a risk of privacy intrusion the following period, expected to cost 7. In the second period, assuming that no intrusion has yet taken place, she may once again protect herself by spending a little more, 6; or she may choose to face a risk of privacy intrusion the next (third) period, expected to cost 9. In the third period she could protect herself for 8 or face an expected cost of 15 in the following last period.

Here too we make no attempt at calibrating the values in Table 2. Again, we focus on the different behavior driven by heterogeneity in time-consistency and sophistication versus naïvete. We assume that $\beta = 1$ for individuals with no self control problems and $\beta = 1/2$ for everybody else. We assume for simplicity that $\delta = 1$ for all.

The time-consistent individuals will obviously choose to protect themselves as soon as possible.

In the first period, naïve time-inconsistent individuals will compare the costs of protecting themselves then or face a privacy intrusion in the second period. Because $5 > 7 * (1/2)$, they will prefer to wait until the following period to protect themselves. But in the second period they will be comparing $6 > 9 * (1/2)$ - and so they will postpone their protection again. They will keep on doing so, facing higher and higher risks. Eventually, they will risk to incur the highest perceived costs of privacy intrusions (note again that we are simply assuming that individuals *believe* there are privacy risks and that they increase over time; we will come back to this concept later on).

Time-inconsistent but sophisticated individuals, on the other side, will adopt a protective technology in period 2 and pay 6. By period 2, in fact, they will (correctly) realize that if they wait till period 3 (which they are tempted to do, because $6 > 9 * (1/2)$), their self-control bias will lead them to postpone adopting the technology once more (because $8 > 15 * (1/2)$). Therefore they predict they would incur the expected cost $15 * (1/2)$, which is larger than 6 - the cost of protecting oneself in period 2. In period 1, however, they correctly predict that they will not wait to protect themselves further than period 2. So they wait till period 2, because $5 > 6 * (1/2)$, at which time they will adopt a protective technology (see also [31]).

To summarize, time-inconsistent people tend not to fully appreciate future risks and, if naïve, also their inability to deal with them. This happens even if they are aware of those risks and they are aware that those risks are *increasing*. As we learnt from the second scenario, time inconsistency can lead individuals to accept higher and higher risks. Individuals may tend to downplay the fact that single actions present low risks, but their repetition forms a huge liability: it is a deceiving aspect of privacy that its value is truly appreciated only after privacy itself is lost. This dynamics captures the essence of privacy and the so-called anonymity sets [32, 14], where each bit of information we reveal can be linked to others, so that the whole is more than the sum of the parts.

In addition, [31] show that when costs are immediate, time-inconsistent individuals tend to procrastinate; when benefits are immediate, they tend to *preoperate*. In our context things are even more interesting because all privacy decisions involve at the same time costs and benefits. So we opt against using eCash [9] in order to save us the costs of switching from credit cards. But we accept the risk that our credit card number on the Internet could be used ma-

[4] One may claim that loyalty cards keep on providing benefits over time. Here we make the simplifying assumption that such benefits are not larger than the future costs incurred after having revealed one's tastes. We also assume that the economy ends in period 4 for all individuals, regardless of when they chose to join the loyalty program.

	Period 1	Period 2	Period 3	Period 4
Protection costs	5	6	8	.
Expected intrusion costs	.	7	9	15

Table 2: (Fictional) costs of protecting privacy and expected costs of privacy intrusions over time.

liciously. And we give away our personal information to supermarkets in order to gain immediate discounts - which will likely turn into price discrimination in due time [3, 26].

We have shown in the second scenario above how sophisticated but time-inconsistent individuals may choose to protect their information only in period 2. Sophisticated people with self-control problems may be at a loss, sometimes even when compared to naïve people with time inconsistency problems (how many privacy advocates do use privacy enhancing technologies all the time?). The reasoning is that sophisticated people are aware of their self-control problems, and rather than ignoring them, they incorporate them into their decision process. This may decrease their own incentive to behave in the optimal way *now.* Sophisticated privacy advocates might realize that protecting themselves from any possible privacy intrusion is unrealistic, and so they may start misbehaving now (and may get used to that, a form of coherent arbitrariness). This is consistent with the results by [36] presented at the ACM EC '01 conference. [36] found that privacy advocates were also willing to reveal personal information in exchange for monetary rewards.

It is also interesting to note that these inconsistencies are not caused by ignorance of existing risks or confusion about available technologies. Individuals in the abstract scenarios we described are aware of their *perceived* risks and costs. However, under certain conditions, the magnitude of those liabilities is almost irrelevant. The individual will take very slowly increasing risks, which become steps towards huge liabilities.

5. DISCUSSION

Applying models of self-control bias and immediate gratification to the study of privacy decision making may offer a new perspective on the ongoing privacy debate. We have shown that a model of rational privacy behavior is unrealistic, while models based on psychological distortions offer a more accurate depiction of the decision process. We have shown why individuals who genuinely would like to protect their privacy may not do so because of psychological distortions well documented in the behavioral economics literature. We have highlighted that these distortions may affect not only naïve individuals but also sophisticated ones. Surprisingly, we have also found that these inconsistencies may occur when individuals perceive the risks from not protecting their privacy as significant.

Additional uncertainties, risk aversion, and varying attitudes towards losses and gains may be confounding elements in our analysis. Empirical validation is necessary to calibrate the effects of different factors.

An empirical analysis may start with the comparison of available data on the adoption rate of privacy technologies that offer immediate refuge from minor but pressing privacy concerns (for example, 'do not call' marketing lists), with data on the adoption of privacy technologies that offer less obviously perceivable protection from more dangerous but also less visible privacy risks (for example, identity theft insurances). However, only an experimental approach over different periods of time in a controlled environment may allow us to disentangle the influence of several factors. Surveys alone cannot suffice, since we have shown why survey-time attitudes will rarely match decision-time actions. An experimental verification is part of our ongoing research agenda.

The psychological distortions we have discussed may be considered in the ongoing debate on how to deal with the privacy problem: industry self-regulation, users' self protection (through technology or other strategies), or government's intervention. The conclusions we have reached suggest that individuals may not be trusted to make decisions in their best interests when it comes to privacy. This does not mean that privacy technologies are ineffective. On the contrary, our results, by aiming at offering a more realistic model of user-behavior, could be of help to technologists in their design of privacy enhancing tools. However, our results also imply that technology alone or awareness alone may not address the heart of the privacy problem. Improved technologies (with lower costs of adoption and protection) and more information about risks and opportunities certainly can help. However, more fundamental human behavioral mechanisms must also be addressed. Self-regulation, even in presence of complete information and awareness, may not be trusted to work for the same reasons. A combination of technology, awareness, *and* regulative policies - calibrated to generate and enforce liabilities and incentives for the appropriate parties - may be needed for privacy-related welfare increase (as in other areas of an economy: see on a related analysis [25]).

Observing that people do not want to pay for privacy or do not care about privacy, therefore, is only a half truth. People may not be able to act as economically rational agents when it comes to personal privacy. And the question whether "do consumers care?" is a different question from "does privacy matter?" Whether from an *economic* standpoint privacy *ought* to be protected or not, is still an open question. It is a question that involves defining specific contexts in which the concept of privacy is being invoked. But the value of privacy eventually goes beyond the realms of economic reasoning and cost benefit analysis, and ends up relating to one's views on society and freedom. Still, even from a purely economic perspective, anecdotal evidence suggest that the costs of privacy (from spam to identity theft, lost sales, intrusions, and the like [30, 12, 17, 33, 26]) are high and increasing.

6. ACKNOWLEDGMENTS

The author gratefully acknowledges Carnegie Mellon University's Berkman Development Fund, that partially supported this research. The author also wishes to thank Jens Grossklags, Charis Kaskiris, and three anonymous referees for their helpful comments.

7. REFERENCES

[1] A. Acquisti, R. Dingledine, and P. Syverson. On the economics of anonymity. In *Financial Cryptography - FC '03*, pages 84–102. Springer Verlag, LNCS 2742, 2003.

[2] A. Acquisti and J. Grossklags. Losses, gains, and hyperbolic discounting: An experimental approach to information security attitudes and behavior. In *2nd Annual Workshop on Economics and Information Security - WEIS '03*, 2003.

[3] A. Acquisti and H. R. Varian. Conditioning prices on purchase history, 2001. Presented at the European Economic Association Conference, Venice, IT, August 2002. http://www.heinz.cmu.edu/~acquisti/papers/privacy.pdf.

[4] G. A. Akerlof. The market for 'lemons:' quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84:488–500, 1970.

[5] G. S. Becker and K. M. Murphy. A theory of rational addiction. *Journal of Political Economy*, 96:675–700, 1988.

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

[7] G. Calzolari and A. Pavan. On the optimality of privacy in sequential contracting. Technical report, Gremaq, University of Toulouse, 2004.

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

[9] D. Chaum. Blind signatures for untraceable payments. In *Advances in Cryptology - Crypto '82*, pages 199–203. Plenum Press, 1983.

[10] R. K. Chellappa and R. Sin. Personalization versus privacy: An empirical examination of the online consumer's dilemma. Forhtcoming, 2004.

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

[12] Community Banker Association of Indiana. Identity fraud expected to triple by 2005, 2001. http://www.cbai.org/Newsletter/December2001/identity_fraud_de2001.htm.

[13] S. Corey. Professional attitudes and actual behavior. *Journal of Educational Psychology*, 28(1):271 – 280, 1937.

[14] C. Diaz, S. Seys, J. Claessens, and B. Preneel. Towards measuring anonymity. In P. Syverson and R. Dingledine, editors, *Privacy Enhancing Technologies - PET '02*. Springer Verlag, 2482, 2002.

[15] ebusinessforum.com. eMarketer: The great online privacy debate, 2000. http://www.ebusinessforum.com/index.asp?doc_id=1785&layout=rich_story.

[16] Federal Trade Commission. Identity theft heads the ftc's top 10 consumer fraud complaints of 2001, 2002. http://www.ftc.gov/opa/2002/01/idtheft.htm.

[17] R. Gellman. Privacy, consumers, and costs - How the lack of privacy costs consumers and why business studies of privacy costs are biased and incomplete, 2002. http://www.epic.org/reports/dmfprivacy.html.

[18] I.-H. Harn, K.-L. Hui, T. S. Lee, and I. P. L. Png. Online information privacy: Measuring the cost-benefit trade-off. In *23rd International Conference on Information Systems*, 2002.

[19] 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.

[20] P. Jehiel and A. Lilico. Smoking today and stopping tomorrow: A limited foresight perspective. Technical report, Department of Economics, UCLA, 2002.

[21] 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.

[22] H. Kunreuther. Causes of underinsurance against natural disasters. *Geneva Papers on Risk and Insurance*, 1984.

[23] D. Laibson. Essays on hyperbolic discounting. MIT, Department of Economics, Ph.D. Dissertation, 1994.

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

[25] G. Lowenstein, T. O'Donoghue, and M. Rabin. Projection bias in predicting future utility. Technical report, Carnegie Mellon University, Cornell University, and University of California, Berkeley, 2003.

[26] A. Odlyzko. Privacy, economics, and price discrimination on the Internet. In *Fifth International Conference on Electronic Commerce*, pages 355–366. ACM, 2003.

[27] T. O'Donoghue and M. Rabin. Choice and procrastination. *Quartely Journal of Economics*, 116:121–160, 2001. The page referenced in the text refers to the 2000 working paper version.

[28] R. A. Posner. An economic theory of privacy. *Regulation*, May-June:19–26, 1978.

[29] R. A. Posner. The economics of privacy. *American Economic Review*, 71(2):405–409, 1981.

[30] Privacy Rights Clearinghouse. Nowhere to turn: Victims speak out on identity theft, 2000. http://www.privacyrights.org/ar/idtheft2000.htm.

[31] M. Rabin and T. O'Donoghue. The economics of immediate gratification. *Journal of Behavioral Decision Making*, 13:233–250, 2000.

[32] A. Serjantov and G. Danezis. Towards an information theoretic metric for anonymity. In P. Syverson and R. Dingledine, editors, *Privacy Enhancing Technologies - PET '02*. Springer Verlag, LNCS 2482, 2002.

[33] A. Shostack. Paying for privacy: Consumers and infrastructures. In *2nd Annual Workshop on Economics and Information Security - WEIS '03*, 2003.

[34] H. A. Simon. *Models of bounded rationality.* The MIT Press, Cambridge, MA, 1982.

[35] P. Slovic. What does it mean to know a cumulative risk? Adolescents' perceptions of short-term and long-term consequences of smoking. *Journal of Behavioral Decision Making*, 13:259–266, 2000.

[36] S. Spiekermann, J. Grossklags, and B. Berendt. E-privacy in 2nd generation e-commerce: Privacy preferences versus actual behavior. In *Proceedings of the ACM Conference on Electronic Commerce (EC '01)*, pages 38–47, 2001.

[37] G. J. Stigler. An introduction to privacy in economics and politics. *Journal of Legal Studies*, 9:623–644, 1980.

[38] P. Syverson. The paradoxical value of privacy. In *2nd Annual Workshop on Economics and Information Security - WEIS '03*, 2003.

[39] C. R. Taylor. Private demands and demands for privacy: Dynamic pricing and the market for customer information. Department of Economics, Duke University, Duke Economics Working Paper 02-02, 2002.

[40] T. Vila, R. Greenstadt, and D. Molnar. Why we can't be bothered to read privacy policies: Models of privacy economics as a lemons market. In *2nd Annual Workshop on Economics and Information Security - WEIS '03*, 2003.

[41] S. Warren and L. Brandeis. The right to privacy. *Harvard Law Review*, 4:193–220, 1890.

[42] N. D. Weinstein. Optimistic biases about personal risks. *Science*, 24:1232–1233, 1989.

[43] A. Whitten and J. D. Tygar. Why Johnny can't encrypt: A usability evaluation of PGP 5.0. In *8th USENIX Security Symposium*, 1999.