The Value of Privacy: Keeping the Money Where the Mouth is

RILE Working Paper 2014-15

Ignacio N. Cofone[*]

Abstract

Despite the high value that internet users place on privacy, they offer their personal information for low compensations. This behavior, known as the privacy paradox, has been explained by a stream of literature with different behavioral biases, and in particular with hyperbolic discounting. However, economics offers two possible reasons to discount payoffs—costs of waiting and hazard rates—which produce two possible reasons for choice reversal. The paper shows that the second can explain the privacy paradox within a rational-choice framework in a way that fits more intuitively with consumer claims and with contemporary policy debates on privacy. This account would change the policy conclusions of the hyperbolic discounting model and would account for current trends in data protection law. In particular, it would be relevant for the right to be forgotten.

JEL Classification: K39; K11; D81; D18

Keywords: online privacy; privacy paradox; right to be forgotten; behavioral advertising; hyperbolic discounting; uncertainty.

[*] Erasmus University Rotterdam, Rotterdam Institute of Law & Economics. Contact: cofone@law.eur.nl. Burgemeester Oudlaan 50, Post box 1738, 3000DR Rotterdam, The Netherlands. +31617684367.

1. Introduction

In agreements that internet users conclude with some websites and service providers—normally called "privacy agreements" or "privacy policies"—, they are allowed to use a certain product and in exchange they allow the company to profit from the personal information collected through the use of that product, which is utilized for targeted advertising. While technically the use of the product is free, from a welfare perspective the price they pay for it is the aspect of their privacy they relinquish.

Social networks are a good example of this. Facebook users, for instance, receive a free account that they use to communicate with their peers or for leisure, and while doing so they reveal personal information about themselves. This information ranges from basic information such as their gender and age up to very personal information such as who they are currently dating, where they travel and what movies they like to watch. With 850 million users that provide this information periodically the company is worth around 100 billion dollars because of the possibilities it presents for behavioral advertisers.

A first approximation to this interaction would suggest that these consumers (hereafter data subjects) will disclose the amount of information for which their marginal cost of disclosure equals their marginal benefit for the use of the product. In such way—one could conclude by applying the first stream of literature in the economics of privacy (Posner 1978; Posner 1981; Stigler 1980; Hirshleifer 1980)—the market would allocate data subjects' personal information to their highest valuer.

With the rise of these products, however, a series of complaints have been made both by consumers and consumer associations that state that their privacy[1] is not being properly protected in the online domain (Norberg, Horne, and Horne 2007). At the same time, privacy-protecting technologies were not a success in the private market. Similarly, people have declared in surveys that they place a very high value on their privacy, while in incentivized experiments they disclose their personal information for little compensation. This behavior is known as the "privacy paradox".

The privacy paradox is a phenomenon that indicates that there is in fact a different incentive structure between internet data collection and processing and the traditional privacy problems long addressed by law. This could justify a different regulation for the former, such as the much-debated right to be

[1] The term "privacy" in this paper represents the right or the ability to keep one's personal information to oneself, or the right to exclude others from one's personal information.

forgotten. As it is clear, however, an answer on which type of regulation is appropriate stems not from the paradox itself but from the explanation of its incentive structure.

The main explanation provided so far for the privacy paradox is that data subjects have cognitive biases (Acquisti, John, and Loewenstein 2013), and in particular they display self-control problems and thus they hyperbolically discount (Acquisti 2004). This paper argues that the privacy paradox, when examined closely, might after all be not an amalgam of behavioral biases but instead an uncertainty problem. The question it asks is then: is it possible to provide an explanation for the privacy paradox within a rational-choice framework?

The next section will review the experimental findings that confirm the existence of a paradox in data subject's behavior. Section 3 will then provide the standard explanation for this paradox together with an alternative explanation. Section 4 will discuss their robustness to explain the paradox. Section 5 will introduce the normative implications of this argument. Section 6 will conclude.

2. The Privacy Paradox

2.1. Inconsistencies in Privacy Valuations

Part of the experimental literature on the topic addresses the question of how much data subjects really value their privacy in view of the mentioned paradox.

The earliest studies in the topic show the inconsistency between data subjects' declared concern for privacy and their actual behavior online (Spiekermann, Grossklags, and Berendt 2001). The study grouped participants according to their own declared privacy concern and discovered that, when participating in an online shopping simulation, there was no significant difference between the groups regarding the amount of personal information revealed.[2]

The same finding can be found in other experiments, which show that privacy concerns announced by subjects prior to the experiment are inconsistent with shopping behavior during the experiment (Berendt, Günther, and Spiekermann 2005). It was also shown that privacy concerns of experimental subjects are a weak predictor of membership and of the amount

[2] This included information such as: in which occasions the subject takes photos, what she does with her pictures, what is her motivation for taking pictures, how photogenic she is, and how conceited she is (Spiekermann, Grossklags, and Berendt 2001).

of information disclosed through social networks (Acquisti and Gross 2006). In an experiment where almost 90% of respondents of a survey declared that they have a high concern about their own privacy, again almost 90% of these respondents accepted to put full name and home address under risk of disclosure in exchange for a loyalty card (Acquisti and Grossklags 2005).

Further experiments have shown that, independently from declared values—which the experiment did not evaluate—data subjects display a low willingness to pay for their personal information. Participants were shown two otherwise identical stores that differed only in the requested information; one store asked for sensitive information and the other for non-sensitive information. When prices between the stores differed all subjects chose to buy from the cheapest store, even if it required more disclosure. More surprisingly, when the prices of the stores were equal, individuals were indifferent between the stores (Beresford, Kübler, and Preibusch 2012).

In more recent studies, data subjects' valuations also display a gap between willingness to pay to protect information and willingness to accept a certain proposal to sell information.[3] In a survey, most participants under a first treatment were not willing to pay one a dollar to prevent behavioral advertising, while under a second treatment most participants were not willing to accept one dollar to allow for behavioral advertising (McDonald and Cranor 2010). In an experiment, subjects were either asked how much money they were willing to pay to protect their otherwise public personal information, or how much they would be willing to accept to allow that information to become public. The average willingness to accept was five times higher than the willingness to pay (WTA:WTP ratio of 5.47) which almost doubles average ratio for other goods (2.92) (Acquisti, John, and Loewenstein 2013).[4]

2.2. Differing Valuations

Other research focuses on whether these valuations vary and whether certain types of personal information are valued more or less than others.

[3] Some studies suggest that, counter intuitively, the offer of a reward for the information actually reduces self-disclosure, intensifying concern against rational-choice based predictions (Andrade, Kaltcheva, and Weitz 2002). Other research, however, contradicts those findings (Hui, Teo, and Lee 2007).

[4] Other related research has explored the impact —or lack thereof— of privacy policies in websites (Andrade, Kaltcheva, and Weitz 2002; Hui, Teo, and Lee 2007), together with other elements such as the mention of a protecting regulation (Spiekermann, Grossklags, and Berendt 2001), privacy seals (Hui, Teo, and Lee 2007), the way in which information requests are presented (Acquisti, John, and Loewenstein 2012), how much other data subjects disclose (Acquisti, John, and Loewenstein 2012) and promises of data breach notifications (Feri, Giannetti, and Jentzsch 2013).

Data subjects value their offline information, which is composed by facts related directly to their person that enters the online domain such as their birth date or health status, differently than their online information, composed by their browsing patterns. In average, they seem to value offline information three times as much as their browsing behavior (Carrascal et al. 2013).

Regarding offline information in particular, it has been shown that—as it is intuitive—data subjects do not value all of its types in the same way. An inverse linear relationship has been shown between the desirability of their personal traits and the value they place on them; people ask for more money in order to reveal their undesirable traits with no direct financial or identity-theft repercussions, such as weight and age (Huberman, Adar, and Fine 2005).

More specifically, some research has indicated that data subjects place different values over different types of offline personal information. For instance, they seem to place a high value on information related to their medical and financial status and information about their families, and to have less trouble disclosing information about product consumption and brand consumption, as well as media usage (Horne and Horne 1998).

This line of research suggests that there are relevant differences between the types of information data subjects choose to disclose—offline and online, sensitive and non-sensitive—, and if personal information is treated as fungible units then experimental results might not be entirely accurate. Since not all types of personal information impact data subjects' utility in the same way, it is questionable to treat personal information as fungible units when analyzing transactions (Wathieu and Friedman 2007).

2.3. Context and Accessibility

Other papers explore the importance of context and accessibility of information at the moment of disclosure, which data subjects seem to respond to. They suggest that consumer behavior is less random than one might think based on the findings reviewed above.

When taking the context of disclosure into account, there is some evidence in favor of agents that behave rationally when facing simple privacy issues (Wathieu and Friedman 2007). Data subjects' privacy concerns seem to be sensitive not only to direct harms—defined as an immediate perceived harm provoked by an information release such as fear of fraud or spam—but also to the indirect consequences of the transmission of information—such as fear of ending up being the object of price discrimination. Data subjects seem mainly concerned about the use that is given to their information—even more than about its transfer (Wathieu and Friedman 2007). An increase in control over the publication of data subjects' personal data decreases their concerns over their privacy and hence increases their willingness to disclose sensitive information (John, Acquisti, and Loewenstein 2011).

Data subject's ability to act in their own self-interest when dealing with privacy issues changes when these issues become complex—there does not seem to be a generalized inability to deal with them (John, Acquisti, and Loewenstein 2011).

Two papers showed that when information about security is made visible, for instance available on browsers themselves, data subjects respond to it (Gideon et al. 2006; Tsai et al. 2011). One of them shows that when information about privacy is available directly on search engines data subjects do prefer websites that offer a higher protection for their privacy, in particular regarding purchases which involve the disclosure of sensitive information (Gideon et al. 2006). The other paper explores whether a different display of privacy policies induces data subjects to incorporate better privacy considerations. It shows that when information is available and salient, data subjects prefer to purchase from retailers that protect their privacy better and are even willing to pay a premium to do so (Tsai et al. 2011).

Finally, while the level of comprehension of privacy policies is very low, and while an accessible link to the privacy policy in websites does not significantly affect levels of disclosure, other more "visceral" notices—such as anthropomorphic elements, self-focused attention mechanisms and a high level of formality in web design—do achieve higher levels of comprehension on data subjects and have an effect on their levels of disclosure (Groom and Calo 2011).

3. Competing Explanations

3.1. Hyperbolic Discounting

An explanation that has been offered for the paradox depicted above is picturing data subjects within a model of self-control problems with agents who hyperbolically discount.[5]

A two-fold explanation has been offered for this. First, if data subjects declare that they value their privacy highly, but then they act disregarding it or offer it for low compensations, this behavior can be interpreted as setting a certain plan of action or consumption pattern (they possess a good which they value highly so they should only sell it for a high price) and then deviating from it (they offer the good for a low compensation). Second, this behavior is

[5] Hyperbolic discounting is an increasing rate of time preference over time so that the distant future is more heavily discounted than the near future.

consistent with the literature that suggests that people in many situations discount the distant future at lower rates than the near future (Thaler 1981).

It has been argued that based on the available data from experimental findings it is unrealistic to expect rationality from data subjects. Behavioral economics has shown individuals display in many occasions hyperbolic discounting, under insurance, self-control problems, and immediate gratification, all of which alter people's ability to make decisions. From this perspective, it was argued that the privacy paradox is an example of these biases driving behavior (Acquisti 2004; Acquisti and Grossklags 2004; Acquisti, John, and Loewenstein 2013).

Based on this line of reasoning, models of self-control problems have been offered as an explanation for the experimental data. Data subjects' behavior has been then explained based on a model of immediate gratification where agents hyperbolically discount when facing decisions involving different delays (O'Donoghue and Rabin 2001). Some of these works have stated that it is unrealistic in many scenarios to assume knowable probabilities or complete beliefs (Acquisti and Grossklags 2007; Acquisti and Grossklags 2005). It has been mentioned, similarly, that data subjects face uncertainty about the possible outcomes, the magnitude of their consequences, the possible measures to protect themselves, the actions taken by those who desire his or her information, and the existence of some unforeseeable events, between others.[6]

Could uncertainty play a more significant role? In economics there are two possible reasons to discount payoffs—costs of waiting and hazard rates—which produce two different reasons for choice reversal. The first, hyperbolic discounting, relies on assumptions on the agents, while the second, decreasing or uncertain hazard rates, relies on assumptions on the context.

The second stage of the explanation given to data subjects' behavior can be put into question. The literature that reconciles non-constant discounting with dynamic consistency can offer a feasible explanation to the observed consumer behavior within rational-choice theory.

3.2. Discounting Based on a Hazard Rate I: Declining Rate

Time discounting means that people care less about future consequences than about present consequences, for any possible reason. A relevant reason to discount is that waiting is costly. Another relevant reason is that, over time, payoffs have the risk of depreciating or disappearing; payoffs have a certain hazard rate.

[6] In addition, these models of discounting abstract from liquidity constrains, and therefore from immediate needs that a subject could have when facing the choice.

Imagine an agent who has two choices. The first one is a choice between a certain payoff *V* (100 Euros) at time *T* (now) or a bigger payoff *V'* (150 Euros) at further time *T'* (a year from now). The second one is between the same payoffs (*V* and *V'*) at times *T+t* (3 months from now) and *T'+t* (a year and 3 months from now) (Sozou 1998).

If the hazard rate that payoffs have is independent of time (in our example, there is a constant chance of the promisor of the 150 euros going bankrupt), then the discount rate of a rational agent should be constant. If the agent has a choice between a payoff *V* at time *T* or a bigger payoff *V'* at further time *T'* with a constant hazard rate (λ), then the expected payoff she compares with *V* should be $e^{-\lambda T'}V'$ (Dasgupta and Maskin 2005). The same applies to the second choice.

On the other hand, if the hazard rate (λ) is not independent but dependent on time (*T*), such that λ=λ*(T)* in a way that λ decreases in *T* ($\lambda' < 0$), then closer payoffs would be discounted at a higher rate than distant payoffs (Dasgupta and Maskin 2005) (each month there is a lower chance of the promisor going bankrupt, because he handles his business better). This increases the discount rate as hyperbolic discounting (Frederick, Loewenstein, and O'Donoghue 2002).

This equivalence between the hyperbolic discounting function and the discount function based on a hazard rate, it should be noted, does not happen with any kind of risk, but only in those cases in which the hazard rate decreases over time. Due to the declining hazard rate the agent is more afraid of the payoff disappearing in the first period, and therefore discounts the first choice and the second choice in a different way. This could happen, for instance, with college students, who have a decreasing hazard rate of dropping out, and start-up firms, which have a decreasing hazard rate of going bankrupt, or the unfortunate promisor of our example.

Using a different example, the following table illustrates this procedure.

Delay	*Probability of survival* (s)	*Hazard rate*	*Expected value of cake* ($s.v_c$)	*Expected value of wine* ($s.v_w$)
No delay	1	0%	2	3
1 month	$\frac{1}{2}$	50%	1	$\frac{3}{2}$
2 months	$\frac{1}{3}$	27%	$\frac{2}{3}$	1
3 months	$\frac{1}{4}$	8%	$\frac{1}{2}$	$\frac{3}{4}$
Intrinsic value of cake v_c =2				

Intrinsic value of wine v_w =3

Table 1. Illustrates choice reversal over a hypothetical choice between cake and wine. A cake now is preferred over wine in one month, but wine in 3 months is preferred over cake in 2 months. Based on the table in Sozou (1998). The hazard rate is added for clarity using the probabilities available in the original table.

In the table, the agent has the choice between cake and wine in different time periods. If she has to choose between cake and wine now, she prefers wine ($2 < 3$). If she has to choose between cake now and wine in a month, she would rather take the cake ($2 > \frac{3}{2}$), since the promise of wine has a probability of 50% to materialize itself in a month ($3 * \frac{1}{2} = \frac{3}{2}$). However, even when fully rational and absent of behavioral biases, if she has to choose between cake in two months and wine in three months, she would rather choose wine ($\frac{1}{2} < \frac{3}{4}$). Her choice between the options shifts because the hazard rate of cake and wine (their marginal probability of disappearing) is decreasing over time. This example illustrates how a rational agent would reverse her choices when placed in a particular context.

3.3. Discounting Based on a Hazard Rate II: Unknown Rate

If we assume that the hazard rate is not declining but it is unknown to the agent, then a rational agent will approach the new situation in a similar way. The first choice she has is only relevant in a conditional way to the payoffs surviving after the first period of time *T* (in our example, now), while her second choice is only relevant conditional to the payoffs surviving after the first period of time *T+t* (in our example, 3 months).

Since the agent faces uncertainty over the hazard rate in each scenario, she will again be afraid of the payoff disappearing in the first period and will use a discount rate that is lower for the second case. So the (rational) agent will behave less patiently in the first choice, even without a declining hazard rate (Sozou 1998; Halevy 2008; Azfar 1999). [7] In the example, if she cannot get the prize that was promised to her at time *T* (now) anyhow, and she must face the unknown risk (she does not know how well the promisor handles his business), she would rather wait longer (*T+t*, 3 months) and get a larger reward.

Rational agents with no strict time preferences that include in their analysis unknown risks such as their own mortality have been shown to display diminishing impatience while maintaining time-consistent choices (Sozou 1998; Azfar 1999), even when the unknown hazard rate is actually

[7] There are also accounts for the change in discounting based on the effect of intervals (Read 2001; Read and Roelofsma 2003)

increasing (Halevy 2005). Non-expected utility models can be the result of rational behavior when faced with uncertainty in the future, which means that preference reversals do not imply impatience when risk is present (Halevy 2008).

This idea has been supported by experimental evidence, where it was found that when uncertainty for the present increased from 0 to 0.5, subjects choosing immediate rewards in standard choice reversal problems decreased from 82% to 39% (present bias is reduced when the present is also risky) (Keren and Roelofsma 1995; Weber and Chapman 2005). A procedure to test for hyperbolic discounting while controlling for uncertainty has also been presented (Fernandez-Villaverde and Mukherji 2006; Besharov and Coffey 2003).[8] Other studies indicate that this is also the case for uncertain delays (McGuire and Kable 2012; McGuire and Kable 2013). This implies incurring in fewer assumptions than an uncertain hazard rate since in such case the prize (or penalty) is certain but only its moment of execution is uncertain.

This is how a rational agent facing two equivalent choices with different delays—even without a declining hazard rate—might have delay-dependent discounting and exhibit choice reversal when the hazard rate is uncertain. A mathematical explanation of this can be found in the appendix of the paper, showing how the uncertainty-based discount function can take the same shape as a hyperbolic discount function.

A way of discriminating between preference-based (dynamically inconsistent) diminishing impatience and uncertainty-based (dynamically consistent) diminishing impatience is checking for preferences towards pre-commitment or flexibility (Casari 2009). While sophisticated agents who discount due to self-control problems should be willing to pay to pre-commit—since that would maximize their long-term utility—sophisticated agents discounting due to uncertainty should be willing to pay for flexibility—since their utility is increased by the ability to adjust to new information.[9]

Agents facing a temptation problem who are aware of that problem would be willing to pay for a mechanism to bind oneself in order to avoid changing a certain decision in the future, and hence resist temptation. Some

[8] Agents in an experiment have consumption choices involving immediate and delayed consumption, and receive shocks in their preferences before each choice. Agents receiving different shocks initially make different decisions regarding those choices, while as the time horizon is moved forward shocks become irrelevant. The demand for pre-commitment devices during the whole experiment is very low.

[9] Sophisticated agents are defined as those who are aware of the reason for the diminishing impatience, while naïve agents are defined as those who are not. Naïve agents, of course, are not willing to pay for either pre-commitment or for flexibility.

common examples of this mechanism are not having alcohol or unhealthy food at home, or taking a limited amount of money—and no credit cards—to the casino. Agents facing an uncertainty problem and who are aware of it, on the other hand, would be willing to pay to be able to adapt to the context once expectations becomes certain. Naive agents, of course, will be willing to pay for neither.

The next section evaluates the feasibility of applying these theories to data subjects. Section 4.1 evaluates whether data subjects' behavior resembles the behavior of other agents who face temptation, and section 4.2 evaluates whether their context resembles a context of a decreasing or uncertain hazard rate.

4. Explaining Data Subject Behavior

4.1. Explaining Consumer Claims

One of the key features of self-control problems is that people recognize a certain behavior in themselves that is not consistent with a certain aim they have. They display a certain willingness to stop that behavior and—if sophisticated—they also recognize that they will probably continue to exhibit it. Pre-committing is then an optimal strategy (O'Donoghue and Rabin 2001).

A model of data subjects who hyperbolically discount should lead one to conclude that an optimal strategy for data subjects is pre-commitment. Policy recommendations would then approach the privacy paradox from a paternalistic or libertarian-paternalistic standpoint (Acquisti 2009). The aim of a regulation that takes this interpretation of user behavior into account would be the provision of tools to pre-commit not to disclose personal information. Alternatively, a regulation aiming to do this can create a system of reward substitution—paying to avoid disclosure, charging to disclose, creating guilt, imposing additional obstacles, etc.

On the other hand, a model of uncertainty-based discounting data subjects such as the one this paper suggests would reverse this idea and lead to the conclusion that policy should provide data subjects with flexibility regarding their choices.

In order to distinguish between the two discounting mechanisms, as it was mentioned, it is relevant to ask: are data subjects and consumer associations demanding the introduction of pre-commitment mechanisms regarding their privacy, or are they asking for something else? In this context, regret does not seem to be a central driver of consumer claims. Even taking into account availability bias, most data subjects would say that they experience regret only on a small proportion of their information sharing. At the same time, at least anecdotal evidence shows surprise in data subjects when they discover how much others can know about their personal information based on what they disclose. Social network users do not typically promise themselves to close their profiles and fail to do so—as many dieters, smokers and people who want to do more exercise do—but they are sometimes shocked on how targeted advertisements display the topics they were recently concerned about.

The demands made by consumer associations typically focus on a lack of transparency in personal information processing. The European Data subjects Organization, for instance, has said through one of its members that "data subjects are sleep-walking in a world without privacy. They do not realize their data is being collected and processed" (Warman 2012). Similarly, the popular objection against targeted advertising is a visceral reaction that qualifies it as "creepy" or "spooky" (Schwartz and Solove 2011; Ur et al. 2012). Why would an increase in the relevance of advertising content with regards to data subjects' interests be qualified in such a way? It fits this characterization to state that the reason is that those data subjects are not aware which companies have information about their interests until they find the advertisements and are surprised.

A simple but illustrative example can be obtained by imputing into Google Trends (which illustrates interest over time of chosen key words as expressed in Google searches) the key words "delete data" against "stop sharing data"—or similar alternatives. The numbers obtained, which in the case of those key words are 99% for "delete data" and 0% for "stop sharing data" for September 2013,[10] seem to indicate that data subjects—or at least data subjects searching on Google—are substantially more concerned, and are getting increasingly concerned, about how to delete their data (flexibility) as opposed to how to stop sharing it (pre-commitment).[11]

The experimental data reviewed in the first section also coincides with this. As it was seen, data subjects' reaction to privacy decisions change depending on the complexity of the decisions (Wathieu and Friedman 2007;

[10] The company does not disclose the absolute search volumes, so numbers are relative, being 100% the maximum number of searches for any of the keywords in the chosen period.

[11] For an analysis of this since 2004 until now see http://www.google.com/trends/explore#q=delete%20data%2C%20%20stop%20sharing%20data%2C%20%20avoid%20uploading%20data&cmpt=q (Last time accessed 25/11/2013).

John, Acquisti, and Loewenstein 2011) and when information about privacy becomes more visible data subjects do respond to it by choosing higher privacy protections (Gideon et al. 2006; Tsai et al. 2011). This behavior is less consistent with a model of irrational data subjects that display self-control problems than with a model of data subjects facing an uncertain decision-making scenario, since irrationality would prevent data subjects from responding to these stimuli.

At a more general level, the fact that data subjects systematically display different valuations for different types of information about them (section 2.2) and that they react rationally to changes in context and accessibility (section 2.3) indicates that they have a rational approach towards their personal information. An agent who faces temptation and discounts hyperbolically should discount independently of context, while the available experimental evidence has shown that data subjects respond to context when such context is visible.

4.2. Their Hazard Rate

The case of internet privacy presents two differences with the traditional experiments on hyperbolic discounting. First, while in the latter agents have a choice between a monetary payoff in the present time and a larger monetary payoff in the future, in the case of internet privacy the tradeoff agents have is between a certain penalty in the present time (not extracting the benefits of disclosing their information) and a larger penalty in the future in the form of a privacy breach.[12]

This risk of privacy breach represents the disutility of any use that can be made with the traded information that is unpleasant to the data subject, and hence she should take it into account as an expected cost when deciding whether to disclose. Therefore, it can take many forms, the most common ones ranging from mostly harmless annoyances such as receiving spam email, to more serious consequences such as public disclosure of embarrassing information, the acquisition of information by medical insurers or future employers that would financially damage the person, identity fraud or identity theft.

[12] There is also the element that data subjects face eventual losses instead of gains as it is the case in most of the hyperbolic discounting literature. This is illustrated in O'Donoghue and Rabin (1999): while gains are preferred now better than later, losses are preferred later better than now. This is notwithstanding the fact that people seem to discount losses with a lower discount rate that the one they use to discount gains (Thaler 1981).

The second difference with decision of whether to disclose personal information online is the existence of uncertainty over the outcome. In this case, the privacy breach, which materializes the hazard rate, has an unknown probability of occurrence a each stage.

In experiments where agents have a choice between different monetary payoffs at different times, they are aware of the size of the payoffs and the probability of the larger payoff occurring. Similarly, when people face daily-life situations in which they hyperbolically discount they are aware this as well. In the classic example where someone faces the choice of whether to eat cake or fruit salad, she already knows the extent to which cake might damage her health. If she chooses it still, one can argue she was discounting hyperbolically because she knew about the expected payoff beforehand—if she had not known the adverse health effects of cake then the decision would not have been based on hyperbolic discounting but on blissful ignorance.

On the other hand, when people decide whether to disclose their personal information, they ignore the probabilities of the bigger penalty (privacy breach) occurring in each period—this is, they ignore the hazard rate.

Every time a company trades a data subject's personal information with another, the data subject faces an increase in the risk of a privacy breach. This ranges from legitimate transfers of information, which can be traded to other companies with a different business model, to illegitimate transfers, such as hacking. A large amount of data-dredging practices, for instance, take place through virus attacks. As the number of databases that possess a certain piece of information increases, the more likely such piece of information is to be subject of a virus attack, keeping other conditions stable.

When companies trade this information, the data subject has her personal data out of her range of control—as her consent is not needed any more to trade it—while its use still has the potential of impacting her welfare negatively. This means that there are externalities in data trading (Laudon 1996; Varian 2002). Securing such trade is the main role of companies who buy and sell aggregated data subjects' personal information acting as intermediaries between data collectors and advertisers (Hagel and Rayport 1997). While these exchanges are only agreed by data collectors and the intermediaries, or by the intermediaries and the advertising companies, data subjects also face expected costs from each of trade. These externalities imply an incentive for companies to overuse information, since they do not face all costs, which is aggravated by the fact that data subjects will often not learn about the over-disclosure and hence have no opportunity to discipline such companies (Swire and Litan, 1998, 8).

Hence, even if at the moment of making the decision of whether to disclose the data subject had full information about expected costs and benefits, she would make such decision based on an uncertain hazard rate because the risk of a privacy breach is not dependent on her behavior alone but also on the subsequent behavior of the companies that acquire her data. This fact leads both to difficulties in making a welfare maximizing decision and, as the previous section showed, to a discount function that should resemble hyperbolic discounting.[13]

Moreover, if behavioral research is to be taken into consideration, then the uncertainty over the hazard rate in the case at hand is not only determined by the externalities which are unknown to the data subject. There are arguments to believe that, in addition to this objective uncertainty, data subjects face subjective uncertainty, which could be produced by limitations in their computational capacities or by high information costs that are rationally not incurred.

Data subjects many times do not comprehend privacy policies and license agreements (Milne and Culnan 2004) and the available privacy protection tools (Acquisti and Grossklags 2005). Protecting privacy on internet properly requires technical skills that very few data subjects possess (Turow 2003). Some data subjects ignore the simplest behaviors to engage to protect their privacy, namely to avoid opening unwanted email (spam), sharing files, downloading from non-secure sites and clicking on pop-up advertisements. More than half of all Americans believe that the mere existence of a privacy policy means that companies cannot trade their data (Turow et al. 2009). Several data subjects disclose their birth date in social networks under privacy configurations that make them visible to any other internet user, despite the

[13] An example of these expected costs materializing that has turned famous due to its magnitude is the Sony scandal of 2011. In that year, data from 100 million users were stolen from the Sony Online Entertainment databases, including name, address, birth date, and in many cases debit and credit card information, which was user later on for different types of fraud (Sony Online Entretainment Press Release 2011). Many of those data subjects did not use the device which uploads information to that database, but their information was in the database nonetheless because it had been replicated or moved. That replication, or that movement, had imposed on them an expected cost in the form of an externality of which they were not aware.

fact that the increases the probability of making them victims of identity fraud[14] and identity theft.[15]

Since in cases of online disclosure of personal information the outcomes of the choice—a large negative payoff—do not occur with certainty, a study where a discount rate or discount factor is inferred only from observed consumption will include in the same category both the discounter for the delay of the penalty and the discounter for its hazard rate. Any perceived risk will therefore alter the observed discounting for delay (Sozou 1998; Halevy 2008).

This suggests that models that incorporate uncertainty-based discounting might explain better the privacy paradox.[16]

5. Normative Implications

5.1. The Right to be Forgotten

As with consumer claims, some of the main contemporary debates over the regulation of privacy seem to go in the direction of providing data subjects with additional flexibility regarding their choices of whether to disclose. A relevant example of this is the right to be forgotten, maybe the most debated feature of the EU General Data Protection Regulation proposal,[17] and recognized by the European Court of Justice in Google v. Spain.[18]

What the right to be forgotten intends to provide data subjects with is a right to request entities that collect or process data to erase from their database any piece of information regarding that data subject, regardless of the source.[19]

[14] In the Netherlands, for example, around 5% of the population was a victim of identity fraud in 2012 alone either by phishing or pharming —which are different methods to use an internet site to obtain personal information. See the "Safety Monitor 2012" (Ministry of Security and Justice and the Central Bureau of Statistics 2013).

[15] Comparing survey data with data from the United States Federal Trade Comission, studies seem to indicate that 73% of people underestimate the actual chances of identity theft (Acquisti and Grossklags 2005).

[16] This model is also in line with other approaches to the privacy paradox which state that, while people value their own privacy, there are other things that they simply value more (Egelman, Felt, and Wagner 2012).

[17] Article 17 of the Proposal for a Regulation of the European Parliament and of the Council on the protection of individuals with regard to the processing of personal data and on the free movement of such data.

[18] *Google Spain SL and Google Inc. v Agencia Española de Protección de Datos (AEPD) and Mario Costeja González*, 2014 E.C.J. C-131/12 (May 13).

[19] In commissioner Reading's words, "if an individual no longer wants his personal data to be processed or stored by a data controller, and if there is no legitimate reason for keeping it, the data should be removed from their system" (Reding 2012, 5).

From the perspective of theory that has been evaluated above (sections 3.3 and 4.2) one can see that the right to be forgotten could be valuable to data subjects from a welfare perspective inasmuch as it provides them with additional flexibility when making decisions of whether to share personal information. Data subjects would be able publish a certain piece of information and later on reverse their decision by requesting its deletion. This possibility, as it was mentioned, would be helpful for them if dealing with an uncertain hazard rate in personal information sharing.[20]

This is linked to some of the main arguments in favor of the right. Namely, that individuals should be able to act and speak freely without the fear that this will lead to dire consequences in the future as the freedom to change is an element of self-development (Blanchette and Johnson 2002; Conely 2010; Werro 2009). Similarly, Article 29 Working Party has expressed that control over one's personal information is a central aspect of information privacy, where control is commonly instrumented by consent (A29WP 2011).

This feature, in addition, appears to be a central aim of the regulation. In the Frequently Asked Questions, for instance, it is stated that "a reinforced 'right to be forgotten' will help people better manage data protection risks online" (European Commission 2012).

The right to be forgotten has been, at the same time, criticized from different angles, mainly due to the limits it imposes for freedom of expression (Rosen 2012). The right, indeed, might be socially costly, and it could be difficult to defend on efficiency terms. What this perspective indicates is that the right has some value also beyond a deontological perspective, and that inasmuch as it increases flexibility it is not unreasonable to consider it in a more limited version in the public debate.

5.2. Increasing Transparency

If one takes the argument that the uncertainty that data subjects face is a relevant problem in their interaction with service-providing companies, then one encounters the question on how to design regulations that can help data subjects to reduce it. While it is hardly possible to state a list of prerequisites that will eliminate user uncertainty, providing marginal increases in flexibility is possible from a regulatory standpoint at lower burdens for providing companies than the right to be forgotten would present.

[20] In the Frequently Asked Questions of the proposed regulation, for instance, it is stated that "a reinforced 'right to be forgotten' will help people better manage data protection risks online" (European Commission 2012).

The experimental literature reviewed points to the relevance of context in the disclosure of information. Evaluating the contextual differences that helped at informing data subjects can, potentially, reduce the uncertainty over the hazard rate.

A policy that might appear attractive at a first glance are publicly run privacy seals, but it has been suggested that they are largely ineffective (Hui, Teo, and Lee 2007), while they are also costly to implement. Besides, there are privacy seals programs that have emerged in the private market, reducing the added value of publicly run programs.

A useful first step would be to require complete privacy policies. An example of this insufficiency is the frequent use of Flash cookies by many websites —with the ability to track despite having been deleted— even though few of them mention it in their privacy policies. Another is the even less transparent re-spawning of traditional cookies by data storage mechanisms which other websites —including whitehouse.gov— use while avoiding to mention to their users (Singel 2013).

A second step is working on the transparency of such documents. The topic of the incorporation of privacy policies is inevitably connected to the limits in computational ability data subjects have. Showing the actual document — as opposed to only a hyperlink to it— and allowing data subjects to accept it only after scrolling it to the bottom can marginally improve their incorporation, like it is done with regular contracts —one cannot sign a contract at the top as a manifestation of agreement, even if one does not actually read it when signing at the bottom. This consideration is in line with the experimental literature on the topic (Tsai et al. 2011). Although the difference in the information absorbed by making data subjects screen a document is probably small, given that the cost of the change is negligible it is a useful requirement to consider.

Still, privacy policies in their current state are costly to read; it has been estimated that an average user reading all privacy policies would spend 201 hours per year (McDonald and Cranor 2008).[21] An improvement on this would be to provide a digested summary at the top of privacy policies stating its most relevant elements. While 70% of people consider privacy policies are difficult to understand (Turow, Feldman, and Meltzer 2005), research indicates that when privacy is offered in a clear and understandable way people do value it (Shostack and Syverson 2004). The experimental literature on visceral notices also suggests this measure should have a relevant impact (Groom and Calo 2011). These summaries can also be set to clarify a list of pre-defined facts that

[21] This would cost $3,534 to the average American internet user (McDonald and Cranor 2008).

are considered relevant to increase transparency, such as whether the company is allowed to trade the user's information with others, which uses it is authorized to give to the information it receives, and whether the information is deleted after the user removes it from the system. At a more general level, privacy policies should be made more readable (Kelley et al. 2010; Kelley et al. 2009).

Articles 7(2) and 11 of the EU regulation proposal work in this direction. The first requires that if consent for processing data is requested in a declaration which also addresses other issues then the request must be made distinguishable. The second demands that data controllers have privacy policies that are transparent and easily accessible and that communications to data subjects regarding the processing of their personal data should be done in an intelligible way, with plain language that is adapted to the data subject. So do articles 14(1) and 14(2), regarding information and the access to data, and 19(2), which states that the right to object must be explicitly and intelligibly manifested to the data subject.

Implementations of this idea have been designed in the form of Facebook nudges to raise awareness about privacy (Wang et al. 2013). The nudges introduce visual cues about the audience to which posts are visible, time delays for posts —as Gmail has as an optional feature— and the feedback about potential negative perception of posts based on cue words. An implementation of this can require a disclaimer which warns data subjects when a certain information disclosure will be completely public, as some companies have already done.

Regarding types of information, while it is not always easy to distinguish between sensitive and non-sensitive information, online and offline information can be easily distinguished, and it seems that they should also be treated differently. It has been shown that data subjects value online and offline information very differently (Carrascal et al. 2013), which reflects the fact that most disutility for data subjects stems from disclosure of offline information since all sensitive information is necessarily offline information. Simultaneously, most behavioral advertising —and hence most profit— is taken from browsing patterns, which is online information.

Finally, it is feasible to nudge companies into nudging data subjects, which can be implemented in a similar way to how governments nudge companies into other behaviors such as taking care of the environment. Focusing on compensation instead of punishment —for instance, via partial tax exemptions— has been successful in the past. This focus can be implemented to this issue with means such as the creation of social prestige: for example, a prize to the companies that take care of their data subjects the best. Guidelines that reward companies who follow them are easier to design and to monitor than the creation of a strict regulation that imposes fines to companies who do not follow it, and they can fit with what we know about the way in which data subjects behave.

6. Conclusion

The paper explained the privacy paradox and provided an economic explanation for it not considered in the privacy literature so far. This explanation bases itself on the fact that internet consumers seem not to have different preferences that standard consumers, but a different scenario.

The inference of self-control problems from observed choice reversals depends crucially on the absence of uncertainty (Fernandez-Villaverde and Mukherji 2006). As it was argued, data subjects who have a choice of whether to disclose personal information encounter a high level of uncertainty regarding the likelihood of eventual future penalties. Hence, an uncertainty-based choice reversal model with stable preferences can explain their behavior.

In turn, policy recommendations which equate non-constant discounting and dynamic inconsistency can produce welfare-decreasing outcomes if agents were in fact dynamically consistent (Azfar 1999). The elimination of future choices or pre-commitment, although optimal for hyperbolically discounting agents, is rarely optimal when more information is expected to arrive in the future (Amador, Werning, and Angeletos 2006). This arrival of future information is likely to be the case in contexts where costs depend on future behavior of other agents and where benefits largely depend on network externalities. Data subjects, when encountered with a non-exceptional change of context, many times simply change their minds.

This framework reverts the policy recommendations that stem from the privacy paradox literature. The policy conclusions of an uncertainty-based model fit consumer claims and policy debates better than alternative theories. Pre-commitment is neither what data subjects and consumer associations demand nor what policymakers consider to implement as additional regulation for new technologies. Increasing flexibility with devices such as the right to be forgotten, and reducing of uncertainty with devices such as the requirement for informed consent, on the other hand, seem to match the demands which are currently present in the public debate.

The question that is left to answer is then how much are data subjects willing to pay for the increase in flexibility, in the form of sacrificing parts of other fundamental rights such as freedom of expression, and in the form of the costs of monitoring that such measures would imply. The answer to this question will determine if a society will prefer mild regulations focusing on transparency, or if it will prefer the possibility to be forgotten.

Acknowledgments

I am grateful to Michael Faure, Luigi Franzoni, Juan José Ganuza, Fernando Gómez Pomar, Klaus Heine, Sharon Oded, Alessio Pacces, Hans-Bernd Schaffer, Roger Van den Bergh, Ann-Sophie Vandenberghe, Louis Visscher and Nicolai Wenzel for helpful comments and discussions, as well as to the participants of the annual conferences of the *European Association of Law and Economics* (Aix-Marseille, 2014), the *Spanish Association of Law and Economics* (Malaga, 2014) and the *German Association of Law and Economics* (Ghent, 2014), and to four anonymous reviewers of the *Workshop on the Economics of Information Security* (2015).

References

A29WP. 2011. *Opinion 15/2011 on the Definition of Consent*. Brussels.

Acquisti, Alessandro. 2004. "Privacy in Electronic Commerce and the Economics of Immediate Gratification." In *Proceedings of the Fifth ACM Conference on Electronic Commerce*, edited by Jack Breese, Joan Feigenbaum, and Margo Seltzer, 21–29. New York: ACM Press.

———. 2009. "Nudging Privacy: The Behavioral Economics of Personal Information." *IEEE Security and Privacy Magazine* 7 (6): 82–85.

Acquisti, Alessandro, and Ralph Gross. 2006. "Imagined Communities: Awareness, Information Sharing, and Privacy on the Facebook." In *6th Workshop on Privacy Enhancing Technologies*, 1–22. Cambridge: Robinson College of Cambridge University.

Acquisti, Alessandro, and Jens Grossklags. 2004. "Privacy Attitudes and Privacy Behavior." In *The Economics of Information Security*, edited by L Jean Camp and Stephen Lewis, Chapter 1. Dordrecht: Kluwer.

———. 2005. "Privacy and Rationality in Individual Decision Making." *IEEE Security and Privacy Magazine* 3 (1): 26–33.

———. 2007. "What Can Behavioral Economics Teach Us About Privacy ?" In *Digital Privacy: Theory, Technologies and Practices*, edited by A. Acquisti, S. Gritzalis, C. Lambrinoudakis, and S. de Capitani di Vimercanti, Chapter 18. Auerbach.

Acquisti, Alessandro, Leslie John, and George Loewenstein. 2012. "The Impact of Relative Standards on the Propensity to Disclose." *Journal of Marketing Research* XLIX: 160–74.

———. 2013. "What Is Privacy Worth?" *The Journal of Legal Studies* 42 (2): 249–74.

Amador, Manuel, Ivan Werning, and George-Marios Angeletos. 2006. "Commitment vs. Flexibility." *Econometrica* 74 (2): 365–96.

Andrade, Eduardo B., Velitchka Kaltcheva, and Barton Weitz. 2002. "Self-Disclosure on the Web: The Impact of Privacy Policy, Reward, and Company Reputation." *Advances in Consumer Research* 29: 350–53.

Azfar, Omar. 1999. "Rationalizing Hyperbolic Discounting." *Journal of Economic Behavior & Organization* 38: 245–52.

Berendt, Bettina, Oliver Günther, and Sarah Spiekermann. 2005. "Privacy in E-Commerce." *Communications of the ACM* 48 (4): 101–6.

Beresford, Alastair R., Dorothea Kübler, and Sören Preibusch. 2012. "Unwillingness to Pay for Privacy: A Field Experiment." *Economics Letters* 117 (1): 25–27.

Besharov, Gregory, and Bentley Coffey. 2003. *Reconsidering the Experimental Evidence for Quasi-Hyperbolic Discounting. Duke Department of Economics Working Paper*. Durham.

Blanchette, Jean-François, and Deborah G. Johnson. 2002. "Data Retention and the Panoptic Society: The Social Benefits of Forgetfulness." *The Information Society* 18 (1): 33–45. doi:10.1080/01972240252818216.

Carrascal, Juan Pablo, Christopher Riederer, Vijay Erramilli, Mauro Cherubini, and Rodrigo de Oliveira. 2013. "Your Browsing Behavior for a Big Mac: Economics of Personal Information Online." In *Proceedings of the 22nd International Conference on World Wide Web*. Geneva.

Casari, Marco. 2009. "Pre-Commitment and Flexibility in a Time Decision Experiment." *Journal of Risk and Uncertainty* 38 (2): 117–41.

Conely, C. 2010. "The Right to Delete." In *AAAI Spring Symposium Series. Intelligent Information Privacy Management*.

Dasgupta, Partha, and Eric Maskin. 2005. "Uncertainty and Hyperbolic Discounting." *The American Economic Review* 95 (4): 1290–99.

Egelman, S., A. P. Felt, and D. Wagner. 2012. *Choice Architecture and Smartphone Privacy: There's a Price for That*. WEIS.

European Commission. 2012. "Data Protection Reform: Frequently Asked Questions." *MEMO/12/41*. http://europa.eu/rapid/press-release_MEMO-12-41_en.htm (Last time accessed on 20/11/2013).

Feri, Francesco, Caterina Giannetti, and Nicola Jentzsch. 2013. *Disclosure of Personal Information Under Risk of Privacy Shocks. University of Bologna Working Paper*. Bologna.

Fernandez-Villaverde, Jesus, and Arijit Mukherji. 2006. *Can We Really Observe Hyperbolic Discounting?* 02-008. *Penn Institute for Economic Research Working Paper No. 02-008*. Philadelphia.

Frederick, Shane, George Loewenstein, and Ted O'Donoghue. 2002. "Time Discounting and Time Preference: A Critical Review." *Journal of Economic Literature* 40 (2): 351–401.

Gideon, Julia, Serge Egelman, Lorrie Cranor, and Alessandro Acquisti. 2006. “Power Strips, Prophylactics , and Privacy, Oh My!” In *Proceedings of the Second Symposium on Usable Privacy and Security*, 133 – 144. New York.

Groom, Victoria, and Ryan Calo. 2011. “Reversing the Privacy Paradox: An Experimental Study.” In *Telecommunications Policy Research Conference*. Schertz.

Hagel, John, and Jeffrey Rayport. 1997. “The Coming Battle for Consumer Information.” *Harvard Business Review* 75 (1): 53–65.

Halevy, Yoram. 2005. *Diminishing Impatience: Disentangling Time Preference from Uncertain Lifetime. University of British Columbia Department of Economics Working Paper 05-17.* Vancouver.

———. 2008. “Strotz Meets Allais: Diminishing Impatience and the Certainty Effect.” *The American Economic Review* 98 (3): 1145–62.

Hirshleifer, Jack. 1980. “Privacy. Its Origin, Function and Future.” *Journal of Legal Studies* 9: 649–66.

Horne, Daniel R., and David A. Horne. 1998. “Domains of Privacy: Toward an Understanding of Underlying Factors.” In *Direct Marketing Educators’ Conference*. San Francisco.

Huberman, Bernardo A., Eytan Adar, and Leslie R. Fine. 2005. “Valuating Privacy.” *IEEE Security and Privacy Magazine* 3: 22–25.

Hui, Kai-Lung, Hock Hai Teo, and Sang-Yong Tom Lee. 2007. “The Value of Privacy Assurance: An Exploratory Field Experiment.” *MIS Quarterly* 31 (1): 19–33.

John, Leslie, Alessandro Acquisti, and George Loewenstein. 2011. “Strangers on a Plane: Context-Dependent Willingness to Divulge Sensitive Information.” *Journal of Consumer Research* 37 (5): 858–73.

Kelley, P. G., J. Bresee, L.F. Cranor, and R. Reeder. 2009. *A “Nutrition Label” for Privacy*. SOUPS.

Kelley, P.G., L.J. Cesca, J. Bresee, and L.F. Cranor. 2010. *Standardizing Privacy Notices: An Online Study of the Nutrition Label Approach*. CHI.

Keren, G, and P Roelofsma. 1995. “Immediacy and Certainty in Intertemporal Choice.” *Organizational Behavior and Human Decision Processes* 63 (3): 287–97.

Laudon, Kenneth C. 1996. "Markets and Privacy." *Communications of the ACM* 39 (9): 92–104.

McDonald, Aleecia, and Lorrie Faith Cranor. 2008. "The Cost of Reading Privacy Policies." *ISJLP* 4: 543.

———. 2010. "Beliefs and Behaviors: Internet Users' Understanding of Behavioral Advertising." In *Telecommunications Policy Research Conference*.

McGuire, Joseph T, and Joseph W Kable. 2012. "Decision Makers Calibrate Behavioral Persistence on the Basis of Time-Interval Experience." *Cognition* 124 (2). Elsevier B.V. 216–26.

———. 2013. "Rational Temporal Predictions Can Underlie Apparent Failures to Delay Gratification." *Psychological Review* 120 (2): 395–410.

Milne, George R., and Mary J. Culnan. 2004. "Strategies for Reducing Online Privacy Risks: Why Consumers Read (or Don't Read) Online Privacy Notices." *Journal of Interactive Marketing* 18 (3): 15–29.

Ministry of Security and Justice and the Central Bureau of Statistics. 2013. "Safety Monitor 2012." http://www.cbs.nl/en-GB/menu/themas/veiligheid-recht/publicaties/artikelen/archief/2013/2013-3912-wm.htm?Languageswitch=on (Last time accessed on 09/10/2013).

Norberg, Patricia a., Daniel R. Horne, and David a. Horne. 2007. "The Privacy Paradox: Personal Information Disclosure Intentions versus Behaviors." *Journal of Consumer Affairs* 41 (1): 100–126.

O'Donoghue, Ted, and Matthew Rabin. 1999. "Doing It Now or Later." *The American Economic Review* 89 (1): 103–24.

———. 2001. "Choice and Procrastination." *Quarterly Journal of Economics* 116 (1): 121–60.

Posner, Richard A. 1978. "The Right of Privacy." *Georgia Law Review* 12 (3): 393–422.

———. 1981. "The Economics of Privacy." *The American Economic Review* 71 (2): 405–9.

Read, Daniel. 2001. "Is Time-Discounting Hyperbolic or Subadditive?" *Journal of Risk and Uncertainty* 23 (1): 5–32.

Read, Daniel, and Peter Roelofsma. 2003. "Subadditive versus Hyperbolic Discounting: A Comparison of Choice and Matching." *Organizational Behavior and Human Decision Processes* 91 (2): 140–53.

Reding, Viviane. 2012. *The EU Data Protection Reform 2012: Making Europe the Standard Setter for Modern Data Protection Rules in the Digital Age. SPEECH/12/26.* Munich. http://europa.eu/rapid/press-release_SPEECH-12-26_en.htm.

Rosen, Jeffrey. 2012. "The Right to Be Forgotten." *Stanford Law Review Online* 64: 88–92. http://www.stanfordlawreview.org/online/privacy-paradox/right-to-be-forgotten?em_x=22.

Schwartz, Paul, and Daniel J Solove. 2011. "PII Problem: Privacy and a New Concept of Personally Identifiable Information, The." *NYU Law Review* 86: 1814–94.

Shostack, Adam, and Paul Syverson. 2004. "What Price Privacy? (and Why Identity Theft Is about Neither Identity nor Theft)." *In J.L. Camp and S. Lewis (eds.), Economics of Information Security,* Kluwer: Chapter 11.

Sony Online Entretainment Press Release. 2011. "Sony Online Entretainment Announces Theft of Data from Its Systems." https://www.soe.com/securityupdate/pressrelease.vm. (Last time accessed on 09/10/2013).

Sozou, Peter D. 1998. "On Hyperbolic Discounting and Uncertain Hazard Rates." *Proceedings of the Royal Society of Biological Sciences* 265 (1409): 2015–20.

Spiekermann, Sarah, Jens Grossklags, and Bettina Berendt. 2001. "E-Privacy in 2nd Generation E-Commerce: Privacy Preferences versus Actual Behavior." In *Proceedings of the Third ACM Conference on Electronic Commerce*, edited by Michael Wellman and Yoav Shoham, 38–47. New York: ACM Press.

Stigler, George J. 1980. "An Introduction to Privacy in Economics and Politics." *Journal of Legal Studies* 9: 623–44.

Thaler, Richard H. 1981. "Some Empirical Evidence on Dynamic Inconsistency." *Economics Letters* 8: 201–7.

Tsai, Janice Y., Serge Egelman, Llorrie Cranor, and Alessandro Acquisti. 2011. "The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study." *Information Systems Research* 22 (2): 254–68.

Turow, Joseph. 2003. “Americans and Online Privacy: The System Is Broken.” In *The University of Pennsylvania Annenberg Public Policy Center Report*. Philadelphia.

Turow, Joseph, L. Feldman, and K. Meltzer. 2005. “Open to Exploitation: American Shoppers Online and Offline.” *The Annenberg Public Policy Center of the University of Pennsylvania Report*.

Turow, Joseph, Jennifer King, Chris Jay Hoofnagle, Amy Bleakley, and Michael Hennessy. 2009. *Americans Reject Tailored Advertising and Three Activities That Enable It. University of Pennsylvania Annenberg School for Communication Working Paper*. Philadelphia.

Ur, Blase, Pedro G Leon, Lorrie Faith Cranor, Richard Shay, Yang Wang, and Pedro Giovanni Leon. 2012. “Smart, Useful, Scary, Creepy: Perceptions of Online Behavioral Advertising Perceptions of Online Behavioral Advertising.” In *SOUPS*. ACM Press.

Varian, Hal R. 2002. “Economic Aspects of Personal Privacy.” *Topics in Regulatory Economics and Policy*, 1–12.

Wang, Yang, Pedro Giovanni Leon, Kevin Scott, Xiaoxuan Chen, Alessandro Acquisti, and Lorrie Faith Cranor. 2013. “Privacy Nudges for Social Media: An Exploratory Facebook Study.” In *Proceedings of the 22nd International Conference on World Wide Web Companion*, 763–70.

Warman, Matt. 2012. “EU Fights ‘Fierce Lobbying’ to Devise Data Privacy Law.” *The Telegraph*. http://www.telegraph.co.uk/technology/internet/9069933/EU-fights-fierce-lobbying-to-devise-data-privacy-law.html (Last time accessed on 09/10/2013).

Wathieu, Luc, and Allan Friedman. 2007. *An Empirical Approach to Understanding Privacy Valuation. Harvard Business School Working Paper 07-075*. Boston.

Weber, Bethany J., and Gretchen B. Chapman. 2005. “The Combined Effects of Risk and Time on Choice: Does Uncertainty Eliminate the Immediacy Effect? Does Delay Eliminate the Certainty Effect?” *Organizational Behavior and Human Decision Processes* 96 (2): 104–18.

Werro, F. 2009. “The Right to Inform v. the Right to Be Forgotten: A Transatlantic Clash.” In *Haftungsbereich Im Dritten Millennium (Liability in the Third Millennium)*, edited by A Colombi Ciacchi, C Godt, P Rott, and LJ Smith, 285–300. Baden: Nomos.

Appendix 1

A Formal Explanation of Discounting Based on Uncertainty

The seminal model (Sozou 1998) can be used to formally show the discounting mechanism explained in the paper.

The expected utility model where the discounting is based on this risk of disappearance of the payoff -as opposed to the cost of waiting- can be illustrated as

$$u(\tau) = u_0 s(\tau) \tag{1}$$

where u is the utility derived from the reward (or penalty), $u(\tau)$ is that utility after waiting time τ and u_0 is that utility at time 0, while $s(\tau)$ is the survival function that translates in the probability of the payoff surviving after the delay τ.

The proposed discounting mechanism can be then represented as

$$u(\tau) = \frac{u_0}{1 + k\tau} \tag{2}$$

where k is a constant whose value symbolizes discounting -the larger k the higher the discounting- and $k > 0$.

Hyperbolic discounting rates are identified with the form $\frac{x}{kt}$ as opposed to the form $\frac{x}{t}$ of regular discount functions.[22] The survival function that corresponds to this (hyperbolic) time preference function is then

$$s(\tau) = \frac{1}{1 + k\tau} \tag{3}$$

Now, say that the hazard rate (λ) is constant but unknown, and it is drawn from a known distribution $f(\lambda)$. The hazard rate is different at each period, but each time it is drawn from the same distribution, so it is neither decreasing nor increasing over time. The survival function will be the aggregation of all those hazard rates at different moments, so we can draw the survival function by direct superposition. If λ can take any value of a certain set N, and the probability of $\lambda_1$is p_1, that of λ_2 is p_2, and so forth until λ_N and p_N, then the probability of surviving until time τ is

$$s(\tau) = p_1 e^{-\tau\lambda} + p_2 e^{-\tau\lambda} + \cdots\ p_N e^{-\tau\lambda} \tag{4}$$

[22] The first function converges into the second as t becomes larger.

then by taking the Laplace transform of $f(\lambda)$ the survival function for the reward after a certain delay τ will be

$$s(\tau) = \int_0^\infty e^{-\tau\lambda} f(\lambda)\,\mathrm{d}\lambda \tag{5}$$

which defines the time-preference function precisely in the same way as hyperbolic discounting. If $f(\lambda) = a$ where $a > 0$ and f is continuous at 0, then the value of the integral should decline hyperbolically as τ increases (Azfar 1999).

Now, substituting (3) into (5) we have that $f(\lambda)$ has to satisfy

$$\int_0^\infty e^{-\tau\lambda} f(\lambda)\,\mathrm{d}\lambda = \frac{1}{1 + k\tau} \tag{6}$$

which gives that

$$f(\lambda) = \frac{1}{k} e^{-\lambda/k} \tag{7}$$

Appendix 2

Experimental Design

Based on the existing experimental literature on how to distinguish between dynamically consistent diminishing impatience and dynamically inconsistent diminishing impatience (Casari 2009), an experimental design can be suggested to distinguish uncertainty-based discounting from temptation-based discounting in data subjects.

Subjects of the experiment face a series of decisions which aim to measure their preference either for pre-commitment or for flexibility. Instead of facing a choice between two alternative payments—as in traditional experiments for choice switch—each subject faces two decisions at different times: first, one at in the lab during the experimental session, and second, one over email later on. Since subject do not need to interact, this can be done in a field experiment.

On the baseline treatment (treatment 1), each subject faces a series of choices between a smaller payment in the form of a discount in a shopping simulation accompanied by a privacy loss (SP),[23] and a larger payment in the form of a lower discount with the avoidance of such privacy loss (LP).[24] During the experimental session, subjects can either pre-commit to LP or postpone the choice between SP and LP to the future, deciding upon the reception of an email that will be sent to them after an X amount of days (e.g. one week), when the payment will take place. The date of realization of the payments remains the same regardless of the choice made during the experimental session. Two variants of this choice are then introduced under separate treatments.

Treatment 2 makes pre-commitment costly. In order to do this, it incorporates a lower payment amount for the choice made during the session (with LP'<LP). Under this treatment, subjects choose between LP' and having the later choice between LP and SP, so pre-commitment is not available for free but for a cost in order to measure whether subjects are willing to pay for pre-commitment (WTPC).

Treatment 3 makes flexibility costly. In order to do this, instead of paying a cost to restrict the choice set, it makes subjects pay a cost to make it wider (with LP"=LP'<LP). Under this treatment, subjects choose between LP

[23] E.g., a $10 discount card with the disclosure of the name and email of the person in a website.

[24] E.g., a $7 discount card without the disclosure.

and having the later choice between LP" and SP, so flexibility is only available at a cost, measuring whether subjects are willing to pay for flexibility (WTPF).

The difference between the number of subjects who are willing to decide now under treatment 2 compared to the number of subjects who are willing to do so on the baseline treatment shows how much subjects value pre-commitment in the choice. In turn, the difference between the number of subjects who are willing to decide later under treatment 3 compared to the number of subjects who are willing to do so on the baseline treatment shows how much subjects value flexibility in the choice

Moreover, if the number of subjects who are willing to pay under treatment 2 is higher than the number of subjects who are willing to pay under treatment 3 (WTPC>WTPF), then subjects have a preference for pre-commitment over flexibility. If the reverse is true (WTPC<WTPF), then subjects have a preference for flexibility over pre-commitment. The first shows that they discount dominantly based on temptation, while the second shows that they discount dominantly based on uncertainty.

The conjecture about the results of the experiment is illustrated in figure 1.

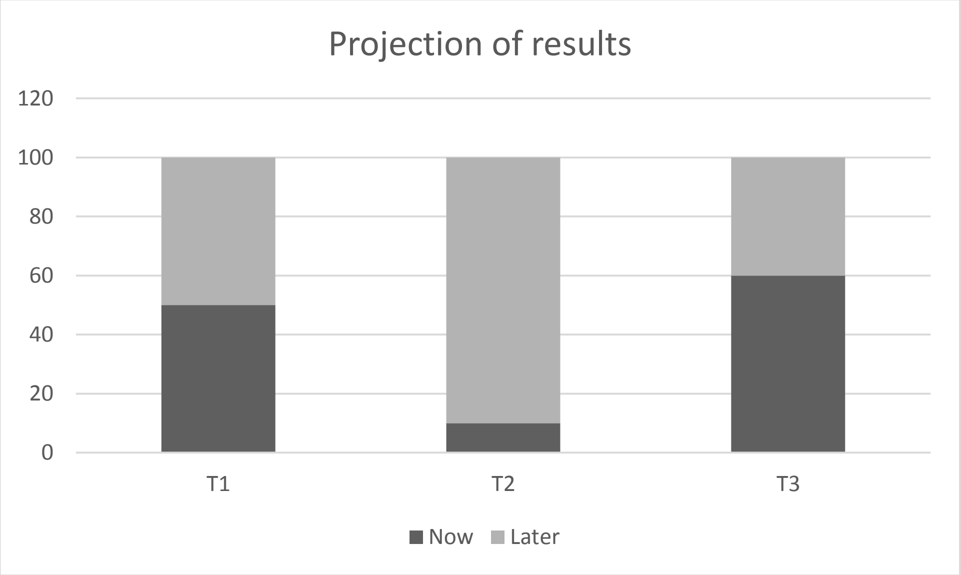

Figure 1. Illustrates expected results of the experiment proposed. Subjects who must pay for pre-commitment prefer to choose later with a significant deviation from the baseline. Subjects who must pay for flexibility prefer to choose now with a slight deviation from the baseline.