What Can Behavioral Economics Teach Us About Privacy?

Alessandro Acquisti[1] and Jens Grossklags[2]

[1] Carnegie Mellon University, Pittsburgh, PA, 15213, USA,
acquisti@andrew.cmu.edu
[2] UC Berkeley, Berkeley, CA, 94720, USA
jensg@sims.berkeley.edu

Draft, preliminary version.
Presented as Keynote Paper at ETRICS 2006.
To appear in:
Digital Privacy: Theory, Technologies and Practices
(Taylor and Francis Group, 2007)

Abstract. Privacy decision making can be surprising or even appear contradictory: we feel entitled to protection of information about ourselves that we do not control, yet willingly trade away the same information for small rewards; we worry about privacy invasions of little significance, yet overlook those that may cause significant damages. Dichotomies between attitudes and behaviors, inconsistencies in discounting future costs or rewards, and other systematic behavioral biases have long been studied in the psychology and behavioral economics literatures. In this paper we draw from those literatures to discuss the role of uncertainty, ambiguity, and behavioral biases in privacy decision making.

1 Introduction

Privacy is a complex decision problem resulting in opinions, attitudes, and behaviors that differ substantially from one individual to another [1]. Subjective perceptions of threats and potential damages, psychological needs, and actual personal economic returns all play a role in affecting our decisions to protect or to share personal information. Thus inconsistencies or even contradictions emerge in individual behavior: sometimes we feel entitled to protection of information about ourselves that we do not control, and end up trading away that same information for small rewards. Sometimes we worry about personal intrusions of little significance, but overlook those that may cause significant damages. In previous works [2–4, 1] we have highlighted a number of difficulties that distance individual actual privacy decision making from that prescribed by classical rational choice theory.[3] First, privacy choices are affected by *incomplete*

[3] According to a straw version of the classical view, individuals would be maximizing their utility over time, using all available information, bayesian updating, and consistent preferences.

information and in particular *asymmetric information* [5]: data subjects often know less than data holders about the magnitude of data collection and use of (un)willingly or (un)knowingly shared or collected personal data; they also know little about associated consequences. Second, the complex life-cycle of personal data in modern information societies can result in a multitude of consequences that individuals are hardly able to consider in their entirety (as human beings, because of our innate *bounded rationality* [6], we often replace rational decision making methods with simplified mental models and heuristics). Third, even with access to complete information and cognitive power to process it exhaustively, various behavioral anomalies and biases could lead individuals to take actions that are systematically different from those predicted by rational choice theory [7]. In this paper, we present an overview of those difficulties, and highlight how research on behavioral economics may improve our understanding of individuals' everyday privacy behavior. In Section 2, we consider the role of asymmetric and incomplete information in privacy scenarios, and how information asymmetries determine risk, uncertainty, and ambiguity in decision making. We argue that due to the prevalence of these informational complications, individuals' privacy relevant behavior may be best understood in terms of bounded rationality [6], and behavioral biases. Specifically, in Section 3 we discuss how insights from the behavioral economic literature may cast a light on the often confusing observations drawn from privacy decision making. In Section 4 we comment on a number of possible paths that privacy research can follow based on those insights.

2 Privacy and Incomplete Information

The occurrence of incomplete information is relevant to privacy for two reasons.[4] The first and perhaps most obvious reason is inherent to the very concept of privacy: an individual has some control on the level of access that other entities can gain on her personal sphere. For example, a subject's personal data may be concealed from other people's knowledge. Other people will thus rely only on incomplete information when interacting with the subject. This is the interpretation of privacy as "concealment" (of job-relevant skills, valuation for a product, creditworthiness, etc.) that Posner and most subsequent formal economic models have recognized [8].

However, incomplete information relates to privacy also in a second sense. It affects the data subject whenever her control on her personal or informational sphere is limited or not clearly determinable. For example, information asymmetries often prevent a subject from knowing when another entity has gained access to or used her personal information; in addition, the subject may not be aware of the potential personal consequences of such intrusions. The associated difficulties to exercise adequate control over private information have been amplified in highly networked, digitized, and interconnected information societies. The release and exchange of personal information has become ubiquitous and

[4] This section is based on "Uncertainty, Ambiguity and Privacy," Alessandro Acquisti and Jens Grossklags, presented at the WEIS 2005 workshop.

often invisible. For example, Varian noted that an individual has little or no control on the secondary use of her personal information, and hence may be subject to externalities whenever other parties transact her personal data [9].

This second sense in which incomplete information creates uncertainties relating to privacy is not new in the economic or legal literature on privacy. However, links between that literature and economic research on incomplete information have been surprisingly limited. So have also been formal or empirical analyses of the impact of risk, uncertainty, or ambiguity on privacy decision making.

Incomplete information complicates privacy decision making because of the resulting mathematical complexity of (evaluating) privacy costs and benefits of transactions. For example, individuals would have to consider multiple layers of outcomes and associated probabilities rather than purely deterministic outcomes. The complexity of the privacy decision environment leads individuals to arrive at highly imprecise estimates of the likelihood and consequences of adverse events, and altogether ignore privacy threats and modes of protection [1].

In the following subsections we restrict the discussion to the role of incomplete information about outcomes and probabilities associated with those outcomes. In particular, we relate the problem of privacy decision making to the research literature in the field of risk, uncertainty, and ambiguity.[5]

2.1 The classical distinction between risk and uncertainty

The distinction between risk and uncertainty in economics dates back to Knight [10] (although earlier discussions of the relations between risk, uncertainty, and utility may be recognized in Bernoulli [11] and then Menger [12]). Knight proposed to distinguish situations characterized by risk (in which the possible random outcomes of a certain event have known associated probabilities) from those characterized by uncertainty or ambiguity (in which the randomness cannot be expressed in terms of mathematical probabilities, and the probabilities themselves are *unknown*). For example, the expected utility theory [13] is based on objectively knowable probabilities (what Knight would have referred to as *risk*).

This distinction has not gone unchallenged by economic theorists and statisticians. A large body of literature suggests that individuals are always able to assign reasonable probabilities to random events. These probabilities could objectively exists in the world [13], and could be used to calculated expected utilities. Or, these probabilities could be *subjective* [14]. Savage adapted expected utility theory into a theory of subjective expected utility, in which, under certain assumptions, people will have personal beliefs about the possible states of nature.[6]

[5] Before we proceed, we want to note that economists, psychologists, and marketers often use terms like risk and uncertainty in different ways. Even within the same discipline researchers disagree on the interpretation given to terms such as *uncertainty*.

[6] The concept of subjective probabilities establishes a bridge between the concept of risk and uncertainty, since the known probability (of a risk) is set on par with a subjective belief. Prescriptively, decision theory and mainstream economic theory of

Behavioral economists and psychologists have worked on modifications of the theories of risk and uncertainty to produce satisfactory descriptive models of human decision making under incomplete information.[7] For example, Hogarth [16] suggests to focus on subjective weights associated to the various possible outcomes of a certain event - where the weights do not have the same mathematical properties as probabilities. In fact, Hogarth proposes that decision weights may be obtained by the individual through a process of anchoring and adjustment. First, an individual may anchor her value on an initial estimate of probabilities over outcomes. Then, she would adjust such an estimate after mentally simulating alternatives values. This adjustment may be influenced by the degree of ambiguity and by the size of the outcome (e.g., whether the gain or loss is large or small).

The debate outlined above is instrumental in the understanding of decision making under uncertainty in both the descriptive and the normative sense. It is also important to the theory of privacy decision making. In particular, we favor the view that in numerous privacy-sensitive situations it is unrealistic to assume existence of known or knowable probabilities or complete (subjective) beliefs for probabilities over all possible outcomes.

2.2 Privacy as a problem of risk or uncertainty?

When presented with a privacy-related problem, consumers often face two major unknowns: a) what privacy-relevant outcomes may occur under different contexts; and b) with what consequences [1]. Implicit in these two major unknowns there are, however, layers of additional uncertainties which we will briefly describe in the following.

First, an individual has often only vague and limited knowledge of the actions she can take to protect (or give away) her personal information. She has also limited knowledge of the possible or actual actions undertaken by other entities (e.g., a marketer's purpose and means to collect information).

Second, actions taken by the individual (whether as an attempt to protect or trade information) or another party have often hardly predictable consequences. For example, it is often unknown whether provided contact information will be used for unwanted communication or whether past consumption data is input for price discrimination strategies.

Third, possible relevant states of nature (with associated additional actions and consequences) may be unknowable in advance, because they depend on

---

expected utility have incorporated the idea that knowledge (or subjective belief) of the actual risks associated with different events and decisions will drive the actions of an economic agent. An economic agent will consider a set of possible actions with different outcomes, probabilities over these outcomes, and associated utilities. He will then choose a strategy consisting of a series of actions leading to the highest expected utility.

[7] Experimental evidence and formal modelling work on ambiguity is reviewed in [15] in great detail.

future, unforeseeable events and environmental changes (e.g., a technology development such as private information retrieval [17]; or Google caching making old Usenet archives searchable).

Fourth, certain desirable actions and information may not be available (see research on asymmetric information and hidden action). Most importantly, consumers often cannot regain control over information formerly released to commercial entities or other individuals.[8]

Fifth, we observed in prior work that individuals iteratively uncover additional layers of a privacy choice situation that reveal further actions and outcomes, with their sets of associated (possible) values and (possible) probabilities. For example, in [1] we describe how people change their perception on which parties have access to their credit card transactional data if they are prompted with this topic repeatedly. We show that individuals sometimes ignore both privacy risks and forms of protection, and even when they are aware of them, often miscalculate their probability of occurrence and their numerical outcome in terms of financial magnitude. This carelessness or ignorance might be justifiable if one considers the effort needed to evaluate everyday privacy choices carefully.

Sixth, privacy protection or invasion are often by-products of other (and sometimes unrelated) transactions. The privacy 'good' is often attached to other goods in complex bundles - or, in other words, trade-offs involving privacy are often trade-offs between heterogeneous goods. For example, when an individual purchases a book online (thus saving the time she would have to spend going to the bookstore and paying in cash), she will often reveal her credit card details to the online merchant, which may lead to an increased risk of identity theft. Or, in order to receive a monetary discount from the grocery store, she will reveal her buying patterns by using a loyalty card, which may increase her probability of receiving junk mail or undesired commercial offers.

Comparisons between those different goods are difficult because of their combinatorial aspects, but may be further complicated if the offers are uncertain or ambiguous. The marketing literature has long been interested in scenarios where the underlying values are incommensurate. In particular, Nunes and Park consider how different forms of wealth are difficult to "convert into any meaningful common unit of measurement." For example, they study a promotion that is presented in nonmonetary terms (e.g., an umbrella) [18]. Under these conditions, the marginal value of the nonmonetary, incremental benefits becomes difficult to evaluate for the individual, in relation to the focal product or its price. Note that privacy-related benefits and costs are rarely monetary and often immaterial.

Because of these intertwined layers of complexity, we conclude that an individual who is facing privacy sensitive scenarios may be uncertain about the values of possible outcomes and their probability of occurrence, and that sometimes she may not even be able to form any beliefs about those values and those probabilities. In fact, she may have no knowledge of the possible *outcomes* of a

[8] There are substantial differences between US and EU data protection legislation concerning the legal rights to gain knowledge of, correct, and delete commercial data records about an individual.

certain situation since the states of nature may be unknown or unknowable in advance. As a result individuals may sometimes ignore both privacy risks and forms of protection.

3 Behavioral economics and privacy

Due to the uncertainties, ambiguities, and complexities that characterize privacy choices, individuals are likely influenced by a number of cognitive limitations and behavioral biases that have been discussed in the literature on behavioral economics.

Behavioral economics studies how individual, social, cognitive and emotional biases influence economic decisions. This research is predominantly based on neoclassical models of economic behavior, but aims to integrate rational choice theory with convincing evidence from individual, cognitive, and social psychology. Behavioral economic models often abandon some of the tenets of rational choice theory: that agents possess consistent preferences between alternatives, chose the utility maximizing option, discount future events consistently, and act upon complete information or known probability distributions for all possible events. In fact, behavioral models expand the economic modelling toolkit by addressing many empirical phenomena - such as how our innate *bounded rationality* limits our ability to exhaustively search for the best alternative; how the *framing* of a scenario or a question may influence an individual's reaction to it; how *heuristics* often replace rational searches for the best possible alternative; and how *biases* and other anomalies affect the way we compare alternatives, perceive risks, or discount values over time [19, 7]. In this section we present a number of themes analyzed in the behavioral literature and discuss their relevance to privacy research, either by making reference to current results or by proposing possible paths of research.

3.1 Helping individuals understand risk and deal with bounded rationality

Consumers will often be overwhelmed with the task of identifying possible outcomes related to privacy threats and means of protection. Even more so, they will face difficulties to assign accurate likelihoods to those states. Policy makers often suggest that providing more information to consumers will help them make better decisions and avoid those impediments. Such additional information may be provided by commercials entities (e.g., anti-spyware vendors), by consumer advocacy groups, or by peers.

However, even if individuals had access to complete information, they would often be unable to process and act optimally on large amounts of data. Especially in the presence of complex, ramified consequences associated with the protection or release of personal information, our innate bounded rationality limits our ability to acquire, memorize and process all relevant information, and it makes us rely on simplified mental models, approximate strategies, and heuristics.

Bounded problem solving is usually neither unreasonable nor irrational, and it needs not be inferior to rational utility maximization. However, these strategies replace theoretical quantitative approaches with qualitative evaluations and "aspirational" solutions that stop short of perfect (numerical) optimization. In [1], we found some evidence of simplified mental models of privacy in a survey about individual privacy attitudes and behavior: a number of survey participants combined together security and privacy issues when they reported feeling that their privacy was protected by merchants who offered SSL connections to complete online payments. Similarly, the presence of a privacy policy may be taken by many to represent privacy protection regardless of its content [51]; or a privacy seal may be interpreted as a guarantee of a trustworthy website [50].

Consumer advocates might suggest that providing individuals with clearly phrased advice or pre-processed information (e.g., to avoid a certain product or activity) will help overcome problems of information overload and bounded decision making. Nevertheless, consumers may still use this data in ways which are different from that of expected utility maximization or contradict their own best interest. In a recent study, Good *et al* found evidence that even well-presented notices of dangerous behaviors of computer programs (e.g., spyware) may not always lead individuals to abort installations or to feel regret about completed installations [20]. Similarly, Spiekermann *et al* individuals' behavior in an interactive online shopping episode was not significantly affected as the privacy statement of the website was modified substantially [22]. Through these studies we learn that individuals are influenced by additional factors that add to the complexity of determining risks and uncertainties associated with privacy threats.

3.2 Framing and heuristics

Tversky and Kahneman have shown that the way a problem or question is *framed* affects how individuals respond to it [24]. In [25] we report experimental evidence from a survey study detailing the impact on the willingness to accept or reject a marketer's privacy related offer, when the consequences of the offer are re-framed in uncertain and highly ambiguous terms. Anecdotal evidence also suggests that it is a safer strategy to convince consumers *ex ante* to provide personal information (even in exchange for small benefits or rewards), than to allow for revelation of privacy-intrusive practices after the fact. In addition, Good *et al* describe preliminary experimental results suggesting that potentially unwanted privacy and security practices discussed in a privacy notice written in vague language might be considered less intrusive by consumers compared to more detailed descriptions of possible dangers [21].

Tversky, Kahneman, and others have also highlighted a number of *heuristics* that guide individual decision making more than rational choice processes. In this context, an heuristic is some technique - often simple and efficient - that helps learning or problem-solving. As an example, individuals often *anchor* on a specific valuation of a good or service, and then adjust that valuation when new information becomes known. However, the process of initial anchoring may

be arbitrary [26], and may create persistent bias in the evaluation process [27]. The value that individuals assign to their own personal information may in fact be assigned through anchoring on a focal and possibly arbitrary value: it is very difficult for an individual to "price" her own information - but once a price has been found (perhaps completely arbitrarily, or perhaps by anchoring it to the reward received by a merchant in exchange for that information) it is likely that the consumer's valuation of her own personal data will thereafter orbit around that value.[9]

Other heuristics may also be found in privacy decision making. For example, individuals may tend to discount as improbable those events that are difficult to picture mentally, such as identity theft (the simulation heuristic [30]); or may associate trustworthy behavior with the neat appearance and design of a website (an example for the representativeness heuristic [31]).

One of the most influential theories in this context is prospect theory [32], that provides an interpretation of how individuals evaluate and compare uncertain gains and losses. Kahneman and Tversky showed that individuals' evaluations around losses and gains can be represented as starting from a reference point, with an S-shaped value function passing through that point. Because of this shape, the same variation in absolute value has larger impact as a loss than as a gain. In other words, this representations reveals how individuals tend to be loss averse, by preferring avoiding losses to acquiring gains. An outcome of the theory is the so-called pseudocertainty effect: individuals tend to make risk-averse choices in the presence of positive expected payoffs but risk-seeking choices in the presence of negative expected payoffs. In addition, individuals are often not only risk averse but also *ambiguity averse* [15]. Given the choice between a certain outcome (e.g., \$10) and a lottery over outcomes (e.g., \$0 with 50% likelihood and X with 50% likelihood), individuals prefer the certain choice unless they are offered a premium in the lottery so that the expected value of the lottery is greater than the certain outcome (e.g., X strictly greater than \$20). Furthermore, there is evidence that competence and knowledge affect individuals' choices. People prefer to bet on events they know more about, even when their beliefs are held constant [33].

The role of these effects on privacy decision making is likely to be significant, although by no means clear - since many competing hypotheses can be formulated. Individuals who do not adopt free and readily available privacy technologies to protect their data, or accept small rewards in exchange for providing their information to parties they know little about, may have simple no interest in keep personal information private or may in fact be displaying both ambiguity love (rather than aversion) or little consideration of future risks. Individuals' low *ex ante* valuation of risks could also be due to a lack of faith about the power of protective solutions to noticeable decrease risks.

Related to prospect theory is also the so-called endowment effect, that suggests that individuals value a good more when they already have it in their

[9] See [28] and [29] for survey and experimental evidence on the valuation of personal information.

possession [34]. In the privacy arena, we found preliminary evidence that individuals tend to assign a higher "sell" value to their personal information (the amount they request from others to give them their information) than the "buy" value (the amount they are willing to spend to make sure that the same information is not released to others) [1]. This happens even for pieces of information that would not appear to have financial consequences when released.

3.3 Other systematic biases

Individuals tend to make sometimes paradoxical, surprising, and seemingly contradictory decisions (see, for example, [35] and [36]).

In the privacy arena, a bias that has been object of attention is hyperbolic discounting. Hyperbolic discounting refers to the idea that people do not discount distant and close events in a consistent way. These inconsistencies could lead to phenomena such as addiction and self-control biases [37]. In [3], we present a model of privacy behavior grounded on some of those distortions - in particular, the tendency to trade-off privacy costs and benefits in ways that may be inconsistent with individuals' initial plans leading to damages of the future selves in favor of immediate gratification [38].

Several other deviations from rationality may also affect the way consumers decide whether to protect or to reveal personal information. We develop below a list of topics for ongoing and future research.

Valence effect. The valence effect of prediction refers to the tendency to overestimate the likelihood of favorable events. In the form of a *self-serving bias*, individuals tend to overestimate the likelihood of favorable events happening to them relative to *other* individuals. In preliminary analysis of users of an online social network, Acquisti and Gross found that online social network users believe that providing personal information publicly on social networks could cause privacy problems to other users, although the same respondents are not particularly concerned about *their* own privacy on those networks [39].

Overconfidence. Overconfidence refers to the tendency to be more confident in one's knowledge or abilities than what would be warranted by facts. Examples of overconfidence can be easily found in different arenas, especially in scenarios where probabilities are difficult to predict. In [1], we found evidence of overconfidence in estimating exposure to a number of privacy risks.

Rational ignorance. Ignorance can be considered rational when the cost of learning about a situation enough to inform a rational decision would be higher than the potential benefit one may derive from that decision. Individuals may avoid assessing their privacy risks for similar reasons: for instance, they may disregard reading a data holder's privacy policy as they believe that the time cost associated with inspecting the notice would not be compensated by the expected benefit (for a related model, see [40]).

Status quo bias. It could also be that individuals choose not to look for solutions or alternatives to deal with their personal information because they prefer, on average, for things to stay relatively the same (the so-called *status quo* bias [41]). In a study of online social networks, we found that the vast majority of users do not change their default (and very permeable) privacy settings [39]. Further research in this area could investigate the relative importance of the *status quo* bias compared to individuals' desire to avoid learning and transaction costs involved in changing existing settings.

Reciprocity and fairness. Reciprocity and fairness have been studied in social psychology and economics [42]. This literature considers the innate desire to act fairly in transactions with other individuals, but also to retaliate or reward others' behavior when deemed appropriate. We believe that such social phenomena are also of relevance to privacy behavior. It is well known, for example, that survey takers are more likely to respond to a survey when the interviewer sends in money even before the survey has been completed. Web sites that ask for registration even before providing any service in return, instead, may end up receiving incorrect or no data at all.

Inequity aversion. Related to the aforementioned concept is the idea of inequity aversion.[10] Because of it, individuals reject offers or express discontent with scenarios in which they feel that *others* are unfairly getting better rewards than the individual, or in which they feel that *they* are getting rewards they do not deserve [43]. In the privacy arena, it is possible that individuals are particularly sensitive to privacy invasions of companies when they feel companies are unfairly gaining from the use of their personal data, without offering adequate consideration to the individual. Or, vice versa, users of social networks may find it natural and fair that, in exchange for the free service they offer, hosting websites end up learning and using information about the individual users [39].

4 How to research the privacy phenomenon?

Privacy decision making is subject to several environmental constraints. Because of the complexities and the influence of uncertainty and ambiguity, providing more privacy information, while helpful, may not be always beneficial to the individual - as it may lead to more cognitive costs, heuristics, and biases.

Heuristics are not necessarily risky strategies. They can be good or bad guides to decision making. Similarly, biases and anomalies that affect privacy behavior are not necessarily damaging. Even *ex post*, only few of the consequences of privacy decisions are actually quantifiable; *ex ante*, fewer yet are. Accordingly, economic actors and observers will find it difficult to judge the optimality of a certain privacy related choice in economic terms.

[10] This concept should not be confused with economic inequality, which typically refers to inequality among individuals or entities within a larger group or society.

It follows that one of the contributions that behavioral economics can offer to privacy research is not necessarily a set of lessons to let consumers avoid *all* costly mistakes, but rather a number of tools to better understand privacy decision making and behavior.

One of the main challenges in this process lies in the fact that several layers of difficulties are intertwined in the privacy phenomenon: incomplete information, framing and heuristics, anomalies and biases may all play interdependent roles, yet by no means all of them may always be present. It is hard, both theoretically and empirically, to separate the impact of any of those layers from the others. Understanding privacy decisions therefore requires a delicate balance between two types of studies: those that cover privacy holistically in its richness, and those consisting in controlled analyses of specific aspects. Combing the recent wave of theoretical models (like [44], [45], or [46]), surveys (like [47], [48], or [1]), and experiments (in the lab [4], [28], or [29]; and in the field [39]) with behavioral economics and newer approaches (such as neuroeconomics [49]) may help us cast light on the intricacies and surprising observations with respect to privacy valuations and actions. The above research can improve policy decision making and technology design for end users and data holding entities. The works and directions discussed in this paper point at an exciting research agenda.

References

1. Acquisti, A., Grossklags, J.: Privacy and rationality in decision making. IEEE Security & Privacy **January-February** (2005) 24–30
2. Acquisti, A.: Privacy and security of personal information: Economic incentives and technological solutions. In Camp, J., Lewis, S., eds.: The Economics of Information Security. (2004) Originally presented at the 2002 Workshop on Economics and Information Security (WEIS '02)
3. Acquisti, A.: Privacy in electronic commerce and the economics of immediate gratification. In: Proceedings of the ACM Conference on Electronic Commerce (EC '04). (2004) 21–29
4. Acquisti, A., Grossklags, J.: Losses, gains, and hyperbolic discounting: An experimental approach to information security attitudes and behavior. In Camp, J., Lewis, S., eds.: The Economics of Information Security. (2004) Originally presented at the 2003 Workshop on Economics and Information Security (WEIS '03)
5. Akerlof, G.A.: The market for 'lemons': Quality uncertainty and the market mechanism. Quarterly Journal of Economics **84**(3) (1970) 488–500
6. Simon, H.A.: Models of bounded rationality. MIT Press, Cambridge, MA (1982)
7. Camerer, C., Lowenstein, G.: Behavioral economics: Past, present, future. In: Advances in Behavioral Economics. (2003) 3–51
8. Posner, R.A.: An economic theory of privacy. Regulation **May-June** (1978) 19–26
9. Varian, H.R.: Economic aspects of personal privacy. In: Privacy and Self-Regulation in the Information Age, National Telecommunications and Information Administration (1996)
10. Knight, F.H.: Risk, Uncertainty, and Profit. Hart, Schaffner & Marx; Houghton Mifflin Company, Boston, MA (1921)

11. Bernoulli, D.: Specimen theoriae novae de mensura sortis. Commentarii Academiae Scientiarum Imperialis Petropolitannae (1738) Translated and published as "Exposition of a New Theory on the Measurement of Risk" in Econometrica 22(1): 23-36, January 1954.
12. Menger, C.: Principles of Economics. New York University Press (1981 edition), New York (1871)
13. von Neumann, J., Morgenstern, O.: Theory of Games and Economic Behavior. Princeton University Press (1953 edition), Princeton, NJ (1944)
14. Savage, L.J.: The Foundations of Statistics. Dover (1972 edition), New York (1954)
15. Camerer, C., Weber, M.: Recent developments in modeling preferences: Uncertainty and ambiguity. The Journal of Risk and Uncertainty **5** (1992) 325–370
16. Hogarth, R.M., Kunreuther, H.: Decision making under uncertainty: The effects of role and ambiguity. In Heller, F., ed.: Decision making and leadership, Cambridge, England: Cambridge University Press (1992) 189–212
17. Chor, B., Goldreich, O., Kushilevitz, E., Sudan, M.: Private information retrieval. In: IEEE Symposium on Foundations of Computer Science. (1995) 41–50
18. Nunes, J.C., Park, C.W.: Incommensurate resources: Not just more of the same. Journal of Marketing Research **40** (2003) 26–38
19. Shefrin, H.: Beyond Greed and Fear: Understanding behavioral finance and the psychology of investing. Oxford University Press (2002)
20. Good, N., Grossklags, J., Mulligan, D., Konstan, J.: Noticing Notice: A large-scale experiment on the timing of software license agreements. Proceedings of SIGCHI conference on Human factors in Computing Systems (CHI'07) **forthcoming** (2007)
21. Good, N., Grossklags, J., Thaw, D., Perzanowski, A., Mulligan, D., Konstan, J.: User choices and regret: Understanding users' decision process about consensually acquired spyware. I/S: A Journal of Law and Policy for the Information Society **2** (2) (2006) 283–344
22. Spiekermann, S., Grossklags, J., Berendt, B.: E-privacy in 2nd generation e-commerce: Privacy preferences versus actual behavior. In: Proceedings of the ACM Conference on Electronic Commerce (EC '01). (2001) 38–47
23. Spiekermann, S.: Product context in EC websites: how consumer uncertainty and purchase risk drive navigational needs In: Proceedings of the ACM Conference on Electronic Commerce (EC '04). (2004) 200–207
24. Tversky, A., Kahneman, D.: The framing of decisions and the psychology of choice. Science **211** (1981) 453–458
25. Acquisti, A., Grossklags, J.: Uncertainty, ambiguity and privacy. In: Workshop on Economics and Information Security (WEIS '05). (2005)
26. Ariely, D., Loewenstein, G., Prelec, D.: Coherent arbitrariness: Stable demand curves without stable preferences. Quarterly Journal of Economics **118** (2003) 73–106
27. Tversky, A., Kahneman, D.: Judgment under uncertainty: Heuristics and biases. Science **185** (1974) 1124–1130
28. Hann, I.H., Hui, K.L., Lee, T.S., Png, I.P.L.: Online information privacy: Measuring the cost-benefit trade-off. In: 23rd International Conference on Information Systems. (2002)
29. Huberman, B., Adar, E., Fine, L.R.: Privacy and deviance. Technical report, HP Labs (2004)
30. Kahneman, D., Tversky, A.: The simulation heuristic. In Kahneman, D., Slovic, P., Tversky, A., eds.: Judgment under uncertainty: Heuristics and biases. Cambridge, UK: Cambridge University Press (1982) 201–210

31. Kahneman, D., Tversky, A.: On the psychology of prediction. Psychological Review **80** (1973) 237–251
32. Kahneman, D., Tversky, A.: Prospect theory: an analysis of decision under risk. Econometrica **47**(2) (1979) 263–269
33. Heath, C., Tversky, A.: Preference and belief: Ambiguity and competence in choice under uncertainty. Journal of Risk and Uncertainty **4** (1991) 5–28
34. Thaler, R.: Towards a positive theory of consumer choice. Journal of Economic Behavior and Organization **1** (1980) 39–60
35. Kahneman, D., Tversky, A.: Choices, Values, and Frames. University Press, Cambridge (2000)
36. Ellsberg, D.: Risk, Ambiguity, and Decision. Garland Publishing, New York and London (2001)
37. Lowenstein, G., Prelac, D.: Choices Over Time. New York: Russell Sage Foundation (1992)
38. Rabin, M., O'Donoghue, T.: The economics of immediate gratification. Journal of Behavioral Decision Making **13**(2) (2000) 233–250
39. Acquisti, A., Gross, R.: Imagined communities: Awareness, information sharing, and privacy on the facebook (2006) Carnegie Mellon University.
40. Vila, T., Greenstadt, R., Molnar, D.: 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. W.Samuelson, Zeckhauser, R.J.: Status quo bias in decision making. Journal of Risk and Uncertainty **1** (1988) 7–59
42. Fehr, E., Gächter, S.: Fairness and retaliation: The economics of reciprocity. Journal of Economic Perspectives **14**(3) (2000) 159–181
43. Fehr, E., Schmidt, K.M.: A theory of fairness, competition, and cooperation. The Quarterly Jornal of Economics **114** (1999) 817–868
44. Taylor, C.R.: Private demands and demands for privacy: Dynamic pricing and the market for customer information. Technical report, Duke University, Economics Department (2002)
45. Acquisti, A., Varian, H.R.: Conditioning prices on purchase history. Marketing Science **24**(3) (2005) 1–15
46. Calzolari, G., Pavan, A.: Optimal design of privacy policies. Technical report, Gremaq, University of Toulouse (2001)
47. Ackerman, M., Cranor, L., Reagle, J.: Privacy in e-commerce: Examining user scenarios and privacy preferences. In: Proceedings of the ACM Conference on Electronic Commerce (EC '99). (1999) 1–8
48. Federal Trade Commission: Privacy online: Fair information practices in the electronic marketplace (2000) http://www.ftc.gov/reports/privacy2000/privacy2000.pdf.
49. Camerer, C., Loewenstein, G., Prelec, D.: Neuroeconomics: How neuroscience can inform economics. Journal of Economic Literature **43**(1) (2005) 9–64
50. Moores, T.: Do Consumers Understand the Role of Privacy Seals in E-Commerce? Communications of the ACM **48**(3) (2005) 86–91
51. Turow, J., Hoofnagle, C.J., Mulligan, D., Good, N., Grossklags, J.: The FTC and Consumer Privacy In the Coming Decade Samuelson Law, Technology and Public Policy Clinic report