

Privacy and Rationality in Individual Decision Making

Alessandro Acquisti

Carnegie Mellon University

Jens Grossklags

University of California, Berkeley

Traditional theory suggests consumers should be able to manage their privacy. Yet, empirical and theoretical research suggests that consumers often lack enough information to make privacy-sensitive decisions and, even with sufficient information, are likely to trade off long-term privacy for short-term benefits.

***To appear in: IEEE Security & Privacy,
January/February 2005, pp. 24-30.***

From its early days^{1,2} to more recent incarnations, economic studies of privacy have viewed individuals as rational economic agents who go about deciding how to protect or divulge their personal information. According to that view, individuals are forward lookers, utility maximizers, Bayesian updaters who are fully informed or base their decisions on probabilities coming from known random distributions. (Some recent works^{3,4} contrast myopic and fully rational consumers, but focus on the latter.) This approach also permeates the policy debate, in which many believe not only that individuals and organizations should have the right to manage privacy trade-offs without regulative intervention, but that individuals can, in fact, use that right in their own best interest.

However, although several empirical studies have reported growing privacy concerns across the US population,^{5,6} recent surveys, anecdotal evidence, and experiments⁷⁻¹⁰ have highlighted an apparent dichotomy between privacy attitudes and actual behavior. First, individuals are willing to trade privacy for convenience or bargain the release of personal information in exchange for relatively small rewards. Second, individuals are seldom willing to adopt privacy protective technologies.

Our research combines theoretical and empirical approaches to investigate the drivers and apparent inconsistencies of privacy decision making and behavior. We present the theoretical groundings to critique

the assumption of rationality in privacy decision making. In addition, we present results from an anonymous, online survey in which we started testing the rationality assumption by analyzing individual knowledge, behavior, and psychological deviations from rationality in privacy-sensitive scenarios.

Challenges in Privacy Decision Making

The individual decision process with respect to privacy is affected and hampered by multiple factors. Among those, incomplete information, bounded rationality, and systematic psychological deviations from rationality suggest that the assumption of perfect rationality might not adequately capture the nuances of an individual's privacy-sensitive behavior.¹¹

First, incomplete information affects privacy decision making because of externalities (when third parties share personal information about an individual, they might affect that individual without his being part of the transaction between those parties),¹² information asymmetries (information relevant to the privacy decision process—for example, how personal information will be used—might be known only to a subset of the parties making decisions), risk (most privacy related payoffs are not deterministic), and uncertainties (payoffs might not only be stochastic, but dependent on unknown random distributions). Benefits and costs associated with privacy intrusions and protection are complex, multifaceted, and context-specific. They are frequently bundled with other products and services (for example, a search engine query can prompt the desired result but can also give observers information about the searcher's interests), and they are often recognized only after privacy violations have taken place. They can be monetary but also immaterial and, thus, difficult to quantify.

Second, even if individuals had access to complete information, they would be unable to process and act optimally on vast amounts of data. Especially in the presence of complex, ramified consequences associated with the protection or release of personal information, our innate *bounded rationality*¹³ limits our ability to acquire, memorize and process all relevant information, and it makes us rely on simplified mental models, approximate strategies, and heuristics. These strategies replace theoretical quantitative approaches with qualitative evaluations and “aspirational” solutions that stop short of perfect (numerical) optimization. Bounded problem solving is usually neither unreasonable nor irrational, and it needs not be inferior to rational utility maximization. However, even marginal deviations by several individuals from their optimal strategies can substantially impact the market outcome.¹⁴

Third, even if individuals had access to complete information and could successfully calculate optimization strategies for their privacy-sensitive decisions, they might still deviate from the rational strategy. A vast body of economic and psychological literature has revealed several forms of systematic psychological deviations from rationality that affect individual decision making.¹⁵ For example, in addition to their cognitive and computational bounds, individuals are influenced by motivational limitations and misrepresentations of personal utility. Experiments have shown an idiosyncrasy between losses and gains (in general, losses are weighted heavier than gains of the same absolute value), and documented a diminishing sensitivity for higher absolute

deviations from the *status quo*. Research in psychology also documents how individuals miscalculate their own future preferences or draw inaccurate conclusions from past choices. In addition, individuals often suffer from self-control problems—in particular, the tendency to trade off costs and benefits in ways that damage their future utility in favor of immediate gratification. Individuals' behavior can also be guided by social preferences or norms, such as fairness or altruism. Many of these deviations apply naturally to privacy-sensitive scenarios.¹¹

Any of these factors might influence decision-making behavior inside and outside the privacy domain, although not all factors need to always be present. Empirical evidence of their influence on privacy decision making would not necessarily imply that individuals act recklessly or make choices against their own best interest. It would, however, imply bias and limitations in the individual decision process that we should consider when designing privacy public policy and privacy-enhancing technologies.

The survey

In May 2004, we contacted potential subjects who had shown interest in participating in economic studies at Carnegie Mellon University. We offered participants a lump sum payment of \$16 to fill out an online, anonymous survey about e-commerce preferences and gathered 119 responses. (We used the term “e-commerce preferences” to mitigate self-selection bias from pre-existing privacy beliefs.) The survey contained several questions organized around various categories: demographics, a set of behavioral economic characteristics (such as risk and discounting attitudes), past behavior with respect to protection or release of personal information, knowledge of privacy risks and protection against them, and attitudes toward privacy. (We discuss only a subset of questions in this article; the full survey is available online at www.heinz.cmu.edu/~acquisti/survey/page1.htm.) This survey was the second round of a research project funded by the Berkman Faculty Development Fund. The first round was a pilot survey we conducted in January, and in the third round, forthcoming, we will further investigate this article's findings.

Participants ranged from 19 to 55 years old (with the mean age of 24). Eighty-three percent were US citizens, with the remainder having heterogeneous backgrounds. More than half of our subjects worked full or part time or were unemployed at the time of the survey, although students represented the largest group (41.3 percent). All participants had studied or were studying at a higher education institution. Hence, our population of relatively sophisticated individuals is not an accurate sample of the US population, which makes our results even more surprising.

Most participants had personal and household incomes below \$60,000. Approximately 16.5 percent reported household incomes above that level, including 6.6 percent with an income greater than \$120,000. Most respondents are also frequent computer users (62.0 percent spend more than 20 hours per week) and Internet browsers (69.4 percent spend more than 10 hours per week) and access computers both at home and work (76.0 percent). Our respondents predominantly use computers running Windows (81.6 percent); 9.6 percent primarily use Macintosh and 8.8 percent use Linux or Unix systems.

Attitudes

A large portion of our sample (89.2 percent) reported to be either moderately or very concerned about privacy (see Table 1). Our subjects provided answers compatible with patterns observed in previous surveys. For example, when asked, “Do you think you have enough privacy in today’s society?,” 73.1 percent answered that they did not. And, when asked, “How do you personally value the importance of the following issues for your own life on a day-to-day basis?,” 37.2 percent answered that information privacy policy was “very important”—less than the shares that believed education policy (47.9 percent) and economic policy (38.0 percent) were very important, but more than the shares of people who believed that the threat of terrorism (35.5 percent), environmental policy (22.3 percent), or same-sex marriage (16.5 percent) were very important.

Privacy attitudes appear correlated with income; the lowest personal income group (less than \$15,000 a year) tended to be less concerned about privacy than all other income groups, with a statistically significant difference in the distributions of concerns by income grouping ($\chi^2 = 17.5, p = 0.008$).

Table 1: Survey results regarding privacy attitudes.

	General privacy concern (%)	Data about offline identity (%)	Data about online identity (%)	Data about personal profile (%)	Data about professional profile (%)	Data about sexual and political identity (%)
High concern	53.7	39.6	25.2	0.9	11.9	12.1
Medium concern	35.5	48.3	41.2	16.8	50.8	25.8
Low concern	10.7	12.1	33.6	82.3	37.3	62.1

Table 1 also shows that requests for identifying information (such as the subject’s name or email address) lead to higher concerns than requests for profiling information (such as age; weight; or professional, sexual, and political profiles). When asked for isolated pieces of personal information, subjects were not highly concerned if the information was not connected to their identifiers. Sensitivity to such data collection practices is generally below the reported general level of concern. However, subjects were more sensitive to data bundled into meaningful groups.⁵ A correlation of data from subjects’ offline and online identities caused strong resistance in 58.3 percent of the sample.

We employed *k*-means multivariate clustering techniques to classify subjects according to their privacy attitudes, extracting base variables used for clustering from several questions related to privacy attitudes. Hierarchical clustering (average linkage) outsets the data analysis. We selected the best partitioning using the Calinski-Harabasz criterion.¹⁶ We derived four distinct clusters: privacy fundamentalists with high concern toward all collection categories (26.1 percent), two medium groups with concerns either focused on the accumulation of data belonging to online or offline identity (23.5 percent and 20.2 percent, respectively), and a group with low concerns in all fields (27.7

percent).

Not surprisingly, concerns for privacy were found to be correlated to how important an individual regards privacy to be. However, by contrasting privacy importance and privacy concerns, we found that for those who most regard privacy as important, concerns were not always equally intense: 46.5 percent of those who declared privacy to be very important expressed lower levels of privacy concerns; in fact, almost 15 percent expressed low absolute concern.

A vast majority of respondents (more than 90 percent) very much agrees with the definition of privacy as ownership and control of personal information. However, a significant number of subjects also care about certain aspects of privacy that do not have immediate informational or monetary interpretation, such as privacy as personal dignity (61.2 percent) and freedom to develop (50.4 percent). In fact, only 26.4 percent strongly agreed with a definition of privacy as the “ability to assign monetary values to each flow of personal information.” Our subjects seemed to care for privacy issues even beyond their potential financial implications.

These results paint a picture of multifaceted attitudes. Respondents distinguish types of information bundles and associated risks, discern between the abstract importance of privacy and their personal concerns, and care for privacy also for nonmonetary reasons.

Behavior

We investigated two forms of privacy-related behavior: self-reported adoption of privacy preserving strategies and self-reported past release of personal information.

We investigated the use of several privacy technologies or strategies and found a multifaceted picture. Usage of specific technologies was consistently low—for example, 67.0 percent of our sample never encrypted their emails, 82.3 percent never put a credit alert on their credit report, and 82.7 percent never removed their phone numbers from public directories. However, aggregating, at least 75 percent did adopt at least one strategy or technology, or otherwise took some action to protect their privacy (such as interrupting purchases before entering personal information or providing fake information in forms).

These results indicate a multifaceted behavior: because privacy is a personal concept, not all individuals protect it all the time. Nor do they have the same strategies or motivations. But most do act.

Several questions investigated the subjects’ reported release of various types of personal information (ranging from name and home address to email content, social security numbers, or political views) in different contexts (such as interaction with merchants, raffles, and so forth). For example, 21.8 percent of our sample admitted having revealed their social security numbers for discounts or better services or recommendations, and 28.6 percent gave their phone numbers. A cluster analysis of the relevant variables revealed two groups, one with a substantially higher degree of information revelation and risk exposure along all measured dimensions (64.7 percent) than the other (35.3 percent). We observed the most significant differences between the two clusters in past behavior regarding the release of social security numbers and descriptions of professional occupation, and the least significant differences for name and nonprofessional interests.

When comparing privacy attitudes with reported behavior,

individuals' generic attitudes might often appear to contradict the frequent and voluntary release of personal information in specific situations.⁷⁻¹⁰ However, from a methodological perspective we should investigate how psychological attitudes relate to behavior under the same scenario conditions (or frames), because a person's generic attitude might be affected by different factors than those influencing her conduct in a specific situation.¹⁷ Under more specific frames, we found supporting evidence for an attitude/behavior dichotomy. For example, we compared stated privacy concerns to ownership of supermarket loyalty cards. In our sample 87.5 percent of individuals with high concerns toward the collection of offline identifying information (such as name and address) signed up for a loyalty card using their real identifying information. Furthermore, we asked individuals about specific privacy concerns they have (participants could provide answers in a free text format) and found that of those who were particularly concerned about credit card fraud and identity theft only 25.9 percent used credit alert features. In addition, of those respondents that suggested elsewhere in the survey that privacy should be protected by each individual with the help of technology, 62.5 percent never used encryption, 43.7 percent do not use email filtering technologies, and 50.0 percent do not use shredders for documents to avoid leaking sensitive information.

Analysis

These dichotomies do not imply irrationality or reckless behavior. Individuals make privacy-sensitive decisions based on multiple factors, including (but not limited to) what they know, how much they care, and how costly and effective their actions they believe can be. Although our respondents displayed sophisticated privacy attitudes and a certain level of privacy-consistent behavior, their decision process seems affected by incomplete information, bounded rationality, and systematic psychological deviations from rationality.

Armed with incomplete information

Survey questions about respondents' knowledge of privacy risks and modes of protection (from identity theft and third-party monitoring to privacy-enhancing technologies and legal means for privacy protection) produced nuanced results. The evidence points to an alternation of awareness and unawareness from one scenario to the other^{18,19} (a cluster of 31.9 percent of respondents displayed high unawareness of simple risks across most scenarios).

On the one hand, 83.5 percent of respondents believe that it is most or very likely that information revealed during an e-commerce transaction would be used for marketing purposes; 76.0 percent believe that it is very or quite likely that a third party can monitor some details of usage of a file sharing client; and 26.4 percent believe that it is very or quite likely that personal information will be used to vary prices during future purchases. On the other hand, most of our subjects attributed incorrect values to the likelihood and magnitude of privacy abuses. In a calibration study, we asked subjects several factual questions about values associated with security and privacy scenarios. Participants had to provide a 95-percent confidence interval (that is a low and high estimate so that they are 95 percent certain that the true value will fall within these limits) for specific privacy-related

questions. Most answers greatly under or overestimated the likelihood and consequences of privacy issues. For example, when we compared estimates for the number of people affected by identity theft (specifically for the US in 2003) to data from public sources (such as US Federal Trade Commission), we found that 63.8 percent of our sample set their confidence intervals too narrowly—an indication of overconfidence.²⁰ Of those individuals, 73.1 percent underestimated the risk of becoming a victim of identity theft.

Similarly, although respondents realize the risks associated with links between different pieces of personal data, they are not fully aware of how powerful those links are. For example, when asked, “Imagine that somebody does not know you but knows your date of birth, sex, and zip code. What do you think the probability is that this person can uniquely identify you based on those data?,” 68.6 percent answered that the probability was 50 percent or less (and 45.5 percent of respondents believed that probability to be less than 25 percent). According to Carnegie Mellon University researcher Latanya Sweeney,²¹ 87 percent of the US population may be uniquely identified with a 5-digit zip code, birth date, and sex.

In addition, 87.5 percent of our sample claimed not to know what Echelon (an alleged network of government surveillance) is; 73.1 percent claimed not to know about the FBI’s Carnivore system; and 82.5 percent claimed not to know what the Total Information Awareness program is.

Our sample also showed a lack of knowledge about technological or legal forms of privacy protection. Even in our technologically savvy and educated sample, many respondents could not name or describe an activity or technology to browse the Internet anonymously to prevent others from identifying their IP address (over 70 percent), be warned if a Web site’s privacy policy was incompatible with their privacy preferences (over 75 percent), remain anonymous when completing online payments (over 80 percent), or protect emails so that only the intended recipient can read them (over 65 percent). Fifty-four percent of respondents could not cite or describe any law that influenced or impacted privacy. Respondents also had a fuzzy knowledge of general privacy guidelines. For example, when asked to identify the OECD Fair Information Principles,²² some incorrectly stated that they include litigation against wrongful behavior and remuneration for personal data (34.2 percent and 14.2 percent, respectively).

Bounded rationality

Even if individuals have access to complete information about their privacy risks and modes of protection, they might not be able to process vast amounts of data to formulate a rational privacy-sensitive decision. Human beings’ rationality is bounded, which limits our ability to acquire and then apply information.¹³

First, even individuals who claim to be very concerned about their privacy do not necessarily take steps to become informed about privacy risks when information is available. For example, we observed discrepancies when comparing whether subjects were informed about the policy regarding monitoring activities of employees and students in their organization with their reported level of privacy concern. Only 46 percent of those individuals with high privacy concerns claimed to have informed themselves about the existence and content of an organizational monitoring policy. Similarly, from the group of

respondents with high privacy concerns, 41 percent admit that they rarely read privacy policies.²³

In addition, in an unframed (that is, not specific to privacy) test of bounded rationality, we asked our respondents to play the beauty contest game, which behavioral economists sometimes use to understand individuals' strategizing behavior. (See the sidebar "Bounded rationality and the beauty contest game" for a full description.) While some of our subjects (less than 10 percent) followed the perfectly rational strategy, most seemed to be limited to a few clearly identifiable reasoning steps.

[sidebar 1 here]

This result does not imply bounded rationality in privacy-relevant contexts; it just demonstrates the subjects' difficulties to navigate in complex environments. However, we found evidence of simplified mental models also in specific privacy scenarios. For example, when asked the open-ended question, "You completed a credit-card purchase with an online merchant. Besides you and the merchant Web site, who else has data about parts of your transaction?," 34.5 percent of our sample answered "nobody," 21.9 percent indicated "my credit card company or bank," and 19.3 percent answered "hackers or distributors of spyware." How is it possible that 34.5 percent of our respondents forget to think of their own bank or other financial intermediaries when asked to list which parties would see their credit-card transactions? When cued, obviously most people would include those parties too. Without such cues, however, many respondents did not consider obvious options. The information is somehow known to the respondents but not available to them during the survey—as it might not be at decision-making time in the real world. In other words, the respondents considered a simplified mental model¹³ of credit-card transactions. (We found similar results in questions related to email and browsing monitoring.)

Further evidence of simplified mental models comes from comments that expanded respondents' answers. For example, some commented that if a transaction with the merchant was secure, nobody else would be able to see data about the transaction. However the security of a transaction does not imply its privacy. Yet, security and privacy seem to be synonyms in simplified mental models of certain individuals. Similar misconceptions were found related to the ability to browse anonymously by deleting browser cookies or to send emails that only the intended recipient can open by using free email accounts such as Yahoo mail.

Similarly, a small number of subjects that reported to have joined loyalty programs and to have revealed accurate identifying information also claimed elsewhere in the survey that they had never given away personal information for monetary or other rewards, showing misconceptions about their own behavior and exposure to privacy risks. (We tested for information items commonly asked for during sign-up processes for loyalty programs, such as name [4.2 percent exhibited such misconceptions], address [10.1 percent], and phone number [12.6 percent].)

Psychology and deviations from rationality

Even with access to complete information and unbounded ability to process it, human beings are subject to numerous psychological

deviations from rationality that a vast body of economic and psychological literature has highlighted: from hyperbolic discounting to underinsurance, optimism bias, and others.¹⁵ (In previous works^{11,24} we discussed which deviations are particularly relevant to privacy decision making.) Corroborating those theories with evidence generally requires experimental tests rather than surveys. Here we comment on indirect, preliminary evidence in our data.

We have already discussed overconfidence in risk assessment and misconception about an individual's information exposing behavior.

Discounting might also affect privacy behavior (see the "Time-inconsistent discounting" sidebar for a detailed explanation). If individuals have time inconsistencies of the form we describe, they might easily fall for marketing offers that offer low rewards now and a possibly permanent negative annuity in the future. Moreover, although they might suffer in every future time period from their earlier mistake, they might decide against incurring the immediate cost of adopting a privacy technology (for example, paying for an anonymous browsing service or a credit alert) even when they originally planned to.¹¹ In an unframed test of our sample, 39.6 percent acted time consistently according to the classical economic perception ($\beta = 1$). However, 44.0 percent acted time inconsistently by discounting later periods at a higher rate (16.4 percent could not be assigned to any of these two categories).

[sidebar 2 here]

Although the discounting results we discuss are not framed to privacy behavior, preliminary evidence about the use of protective technologies is compatible with the theory of immediate gratification. The share of users of a privacy-related technology seems to decrease with the length of time before the penalty from privacy intrusion that technology is supposed to protect will be felt. For example, 52.0 percent of our respondents regularly use their answering machine or caller-ID to screen calls, 54.2 percent have registered their number in a do-not-call list, and 37.5 percent have demanded to be removed from specific calling lists (when a marketer calls them). However, as we noted earlier, 82.3 percent have never put a credit alert on their credit report (of those, however, 34.2 percent are not aware of this possibility at all): the negative consequences of not using this kind of protection could be much more damaging than nuisances associated with unwanted phone calls, but are also postponed in time and uncertain, while the activation costs are immediate and certain. (From our calibration study, we know that 17.4 percent of those individuals that did not use the credit alert option of their credit card company even overestimated the risk of becoming a victim of identity theft.) We will subject these findings to further scrutiny to differentiate between alternative explanations that might be valid, such as lack of knowledge or trust in the accuracy of a technology or a service.

Based on theoretical principles and empirical findings, we are working toward developing models of individual's privacy decision-making that recognize the impact of incomplete information, bounded rationality, and various forms of psychological deviations from rationality.

Many factors affect privacy decision making, including personal

attitudes, knowledge of risks and protection, trust in other parties, faith in the ability to protect information, and monetary considerations. Our preliminary data show that privacy attitudes and behavior are complex but are also compatible with the explanation that time inconsistencies in discounting could lead to under-protection and over-release of personal information. In conclusion, we do not support a model of strict rationality to describe individual privacy behavior. We plan further work on understanding and modeling these behavioral alternatives and on their experimental validation.

Even in our preliminary data, we find implications for public policy and technology design. The current public debate on privacy seems anchored on two prominent positions: either consumers should be granted the right to manage their own privacy trade-offs, or the government should step in to protect the consumer. Our observations suggest that several difficulties might obstruct even concerned and motivated individuals in their attempts to protect their privacy.

While respondents' actual knowledge about law and legislative recommendations was weak, many favored governmental legislation and intervention as a means for privacy protection (53.7 percent). Our test population also supported group protection through behavioral norms (30.6 percent) and self-protection through technology (14.9 percent). Nobody favored the absence of any kind of protection; only one subject suggested self-regulation by the private sector. This is a striking result, contrasting the traditional assumption that US citizens are skeptical toward government intervention and favor industry-led solutions.

Acknowledgments

We thank Nicolas Christin, Lorrie Cranor, Rachel Greenstadt, Adam Shostack, Cristina Wong, the participants at the WEIS 2004 Workshop, Blackhat 2004 Conference, CIPLIT 2004 Symposium, CACR 2004 Conference, and seminar participants at Carnegie Mellon University, the Information Technology department at HEC Montréal, and the University of Pittsburgh for many helpful suggestions.

References

1. R.A. Posner, "An Economic Theory of Privacy," *Regulation*, May–June 1978, pp. 19–26.
2. G.J. Stigler, "An Introduction to Privacy in Economics and Politics," *J. Legal Studies*, vol. 9, 1980, pp. 623–644.
3. A. Acquisti and H.R. Varian, "Conditioning Prices on Purchase History," *Marketing Science*, forthcoming, 2005.
4. C.R. Taylor, "Consumer Privacy and the Market for Customer Information," *RAND J. Economics*, Winter 2004, vol. 35, no. 4, pp. 631–651.
5. M. Ackerman, L. Cranor, and J. Reagle. "Privacy in E-Commerce: Examining User Scenarios and Privacy Preferences," *Proc. ACM Conf. Electronic Commerce (EC 99)*, ACM Press, 1999, pp. 1–8.
6. A.F. Westin, *Harris-Equifax Consumer Privacy Survey 1991*, Equifax, 1991. www.privacyexchange.org/iss/surveys/eqfx.execsum.1991.html.
7. Harris Interactive. *Most People Are 'Privacy Pragmatists' Who, While Concerned about Privacy, Will Sometimes Trade It Off for Other*

- Benefits*, 2003.
www.harrisinteractive.com/harris_poll/index.asp?PID=365.
8. R.K. Chellappa and R. Sin, "Personalization Versus Privacy: An Empirical Examination of the Online Consumer's Dilemma," to be published in *Information Technology and Management*, vol. 6, no. 2–3, 2005.
 9. Il-Horn Hann et al., "Online Information Privacy: Measuring the Cost-Benefit Trade-Off," *Proc. 23rd Int'l Conf. Information Systems*, 2002.
www.comp.nus.edu.sg/~ipng/research/privacy_icis.pdf
 10. S. Spiekermann, J. Grossklags, and B. Berendt, "E-Privacy in Second Generation E-Commerce: Privacy Preferences versus Actual Behavior," *Proc. ACM Conf. Electronic Commerce (EC 01)*, ACM Press, 2001, pp. 38–47.
 11. A. Acquisti, "Privacy in Electronic Commerce and the Economics of Immediate Gratification," *Proc. ACM Conf. Electronic Commerce (EC 04)*, ACM Press, 2004, pp. 21–29.
 12. H. Varian, "Economic Aspects of Personal Privacy," *Privacy and Self-Regulation in the Information Age*, US Dept. of Commerce, 1997.
www.ntia.doc.gov/reports/privacy/privacy_rpt.htm.
 13. H.A. Simon, *Models of Bounded Rationality*, MIT Press, 1982.
 14. N. Christin, J. Grossklags, and J. Chuang, "Near Rationality and Competitive Equilibria in Networked Systems," *Proc. SIGCOMM Workshop on Practice and Theory of Incentives in Networked Systems (PINS)*, ACM Press, 2004, pp. 213–219.
 15. D. Kahneman and A. Tversky, *Choices, Values, and Frames*, Cambridge Univ. Press, 2000.
 16. R.B. Calinski and J. Harabasz, "A Dendrite Method for Cluster Analysis," *Comm. in Statistics*, vol. 3, 1974, pp. 1–27.
 17. M. Fishbein and I. Ajzen, *Belief, Attitude, Intention and Behavior: An Introduction to Theory and Research*, Addison-Wesley, 1975.
 18. S. Fox et al., *Trust and Privacy Online: Why Americans Want to Rewrite the Rules*, The Pew Internet & American Life, 2000.
 19. J. Turow, *Americans and Online Privacy: The System is Broken*, Annenberg Public Policy Center Report, 2003.
 20. S. Oskamp, "Overconfidence in Case Study Judgments," *J. Consulting Psychology*, vol. 29, 1965, pp. 261–265.
 21. L. Sweeney, "K-Anonymity: A Model for Protecting Privacy," *Int'l J. Uncertainty, Fuzziness, and Knowledge-Based Systems*, vol. 10, 2002, pp. 557–570.
 22. Organization for Economic Cooperation and Development (OECD), *The Recommendation Concerning Guidelines Governing the Protection of Privacy and Transborder Flows of Personal Data*, 1980.
http://europa.eu.int/comm/internal_market/privacy/instruments/ocdeguideline_en.htm.
 23. T. Vila, R. Greenstadt, and D. Molnar, "Why We Can't be Bothered to Read Privacy Policies: Models of Privacy Economics as a Lemons Market," *The Economics of Information Security*, L.J. Camp and S.

Lewis, eds., Kluwer, 2004, pp. 143–154.

24. A. Acquisti and J. Grossklags, “Losses, Gains, and Hyperbolic Discounting: An Experimental Approach to Information Security Attitudes and Behavior,” *The Economics of Information Security*, L.J. Camp and S. Lewis, eds., Kluwer, 2004, pp. 165–178.

[sidebar 1]

Bounded rationality and the beauty contest game

In the most popular form of the beauty contest game,¹ questionnaires and experiments ask subjects to respond to the following question:

Suppose you are in a room with 10 other people and you all play a game. You write down a number between 0 and 100. The numbers are collected, and the average is calculated. The person who wrote the number closer to two-thirds of the average wins the game. What number would you write?

The beauty contest is dominance solvable through iterated elimination of weakly dominated strategies; it leads to the game's unique equilibrium where everybody chooses zero—this is what a rational agent would do if it believed that all other agents are rational.

1. R. Nagel, “Unraveling in Guessing Games: An Experimental Study,” *American Economic Review*, vol. 85, no. 5, Dec. 1995, pp. 1313–1326.

[sidebar 2]

Time-inconsistent discounting

Traditionally, economists model people as discounting future utilities exponentially, yielding the intertemporal utility function where payoffs in later periods, t , are discounted by δ^t , with δ being a constant discount factor. Time-inconsistent (hyperbolic) discounting¹ suggests instead that people have a systematic bias to overrate the present over the future. This notion is captured with a parameter $\beta < 1$ that discounts later periods in addition to the δ . Intuitively, an individual with a $\beta < 1$ will propose to act in the future in a certain way (“I will work on my paper this weekend.”) but, when the date arrives, might change her mind (“I can start working on my paper on Monday.”).

1. T. O’Donoghue and M. Rabin, “The Economics of Immediate Gratification,” *J. Behavioral Decision Making*, vol. 13, 2000, pp. 233–250.