

Pre-publication version

Nudging Privacy

The Behavioral Economics of Personal Information

In 356 B.C., a man started a fire that destroyed the temple of Artemis at Ephesus—one of the seven wonders of the ancient world. Captured by the citizens of the town and sentenced to death, he boasted that the arson had been motivated by the desire to

revelation and protection, with satisfaction shared among the individual, organizations, and society as a whole. Research in this area isn't new. The first explicit economic approaches to privacy appeared near the end of the 1970s, in particular via pioneering work from Chicago School economists such as Richard Posner² and George Stigler.³ Research continued through the 1990s, thanks to contributions from economists keenly interested in information technologies (such as Hal Varian⁴ or Eli Noam⁵) and researchers interested in the privacy “calculus” (such as Robert Laufer, Maxine Wolfe,⁶ Pamela Armstrong, or Mary Culnan⁷). Theoretical and empirical studies in this area burgeoned after 2000 with the advent of the commercial Internet and the explosion of technologies for data dissemination, gathering, and analysis.

However, the same technological advances that, in recent years, have vastly expanded information sharing and mining capabilities, have also made our privacy trade-offs more difficult to navigate, exposing surprising dichotomies between our privacy attitudes and our actual behavior. In words, we've reacted to the new technological panopticon by demanding more privacy (surveys keep finding that privacy is one of American consumers' largest concerns; see, for instance, the Consumer Union 2008 Consumer Reports Poll: “Americans Extremely Concerned

ALESSANDRO
ACQUISTI
Carnegie
Mellon
University

gain fame and immortality.

Today, like 2,000 years ago, many seek notoriety at the price of embarrassment, a tarnished reputation, or even infamy. In 2007, a new Facebook group came under media attention: 30 Reasons Girls Should Call It a Night counted “nearly 150,000 members and a collection of nearly 5,000 photos of young women passed out on the pavement, collapsed in shrubbery, peeing in bushes, and vomiting in toilets (or on themselves).”¹ Most of the subjects had uploaded the photos themselves.

What is it that pushes us to seek fame by misconduct or publicity by sharing embarrassing information with strangers? How do we reconcile these desires with the apparent need for privacy that surveys keep finding so widespread among the American population? In short, what drives individuals to reveal, and to hide, information about themselves to and from others?

The Privacy Trade-Off

Privacy decisions often involve attempting to *control* information flows in order to balance competing interests—the costs and

benefits of sharing or hiding personal information. As such, they're a natural field of study for economics. But traditional economic models have made overly restrictive assumptions about the stability and nature of individual privacy preferences, disregarding the psychological and emotional components of (more or less deliberate) decisions about personal data. Newer approaches, drawing on research in behavioral economics and psychology, offer complementary and richer tools to understand privacy decision making and promising initial results. They might be able to reconcile the human need for publicity with our ostensible desire for privacy.

Broadly speaking, privacy economics deals with informational trade-offs: it tries to understand, and sometimes quantify, the costs and benefits that data subjects (as well as potential data holders) bear or enjoy when their personal information is either protected or shared. Privacy economics also tries to understand how to use market mechanisms, technology, or policy to achieve a desirable balance between information

About Internet Privacy,” (see www.consumersunion.org/pub/core_telecom_and_utilities/006189.html). In actions, however, we seem resigned and almost unfazed in the face of escalating intrusions into our personal data. Even the risk of identity theft (an issue of allegedly great importance to many individuals) seems not to generate significant consumer reaction: of all individuals whose data had been obtained by criminals following the Choicepoint data breach, fewer than 10 percent ever called the company to take advantage of the credit protection, insurance, and monitoring tools Choicepoint made freely available (see www.networkworld.com/news/2007/041007-choicepoint-victim-offers.html).

Some social scientists have implicitly or explicitly assumed that people have stable preferences for privacy, and based on those make sensible, coherent trade-offs between privacy and other desired goals—such as participating or not in online social networks.^{2,3} However, substantial literature in behavioral decision research and behavioral economics documents systematic inconsistencies in consumer choices.⁸ This research shows that preferences are often labile and influenced by contextual factors.⁹ For example, preferences depend on how they're elicited or how choice alternatives are framed, as well as how salient the information available to customers is and what types of comparisons evokes. Given that privacy's tangible and intangible consequences are often difficult to estimate, numerous and subtle effects can likely influence and distort the way we value data protection and act on our concerns. This would determine the likely emergence of behavioral inconsistencies and malleable preferences, as well as the fluctuation of privacy concerns over time.

Privacy and Behavioral Economics

To make sense of these inconsistencies, in 2004 I started applying theories and methodologies from behavioral economics and behavioral decision research to investigate privacy decision making.⁶ At Carnegie Mellon University, we started a research that focused on the cognitive and behavioral biases (from risk aversion to immediate gratification) that hamper behavior in this area. Highlighting privacy preferences' malleability, however, doesn't imply that people make “irrational” or wrong decisions about their information. More subtly, systematic inconsistencies and biases suggest that we need richer theories to understand how challenges and hurdles affect the way we make decisions about our personal information. Such hurdles might stem from a combination of factors: inconsistent preferences and frames of judgment; opposing or contradictory needs (such as the need for publicity combined with the need for privacy); incomplete information about risks, consequences, or

solutions inherent to provisioning (or protecting) personal information; bounded cognitive abilities that limit our ability to consider or reflect on the consequences of privacy-relevant actions; and various systematic (and therefore predictable) deviations from the abstractly rational decision process.

Some of these deviations in the privacy domain might be similar to biases behavioral decision researchers have identified in the consumer choice domain. Others might be peculiar to privacy choices. In the course of various studies, colleagues and I have found, for instance, that individuals are less likely to provide personal information to professional-looking sites than unprofessional ones, or when they receive strong assurances that their data will be kept confidential (see <http://ssrn.com/abstract=1430482>). We've found that individuals assign radically different values to their personal information depending on whether they're focusing on protecting data from exposure or selling away data that would be otherwise protected.¹¹ We've found that they

might also suffer from an illusion of control bias that make them unable to distinguish publication control from control of *access* to personal information.¹²

This is an area of research ripe for further investigation, where contributions from different fields (economics, behavioral decision research, psychology, usability, human-computer interaction, and so forth) can fruitfully come together.^{13–15} Two of its most exciting directions focus on understanding how to reconcile our need for privacy with our need for publicity, and how to turn results about cognitive and behavioral biases in privacy and security decision making into normative design research—something that helps us build better privacy technologies and information policies.

The Soft Paternalism Solution

Behavioral economists' recent focus on "soft" or asymmetric paternalism^{16,17} might offer promising tools in this regard. The idea behind soft paternalism is to design systems so that they enhance (and sometimes influence) individual choice to increase individual and societal welfare. To do so, behavioral economists might even design systems to "nudge" individuals, sometimes exploiting the very fallacies and biases they uncover, turning them around in ways that don't diminish users' freedom but offer them the option of more informed choices. Hence, nudging privacy—that is, using soft paternalism to address and improve security and privacy decisions—might be an appealing concept for policy makers and technology designers. This concept goes beyond concurrent attempts at making our computer systems more "usable." Consider, for instance, online social network users who post their dates of birth online. This information isn't particularly sensitive per se, but

could lead to inferences of sensitive data (such as the individuals' Social Security numbers, as our research has shown).¹⁸ A strong paternalistic approach would ban the public provision of birth dates in online profiles—certainly too gross a measure, given that users might have very legitimate reasons to make that information available to others, and that the risks of adverse effects for specific users are limited. A "usability" approach would design a system that makes it easy or intuitive for users to change the visibility settings for their birth dates. A soft paternalistic approach might, instead, provide context to aid the user's decision—such as visually representing how many other users (or types of users) might be able to access that information or what they can do with it; or, it might alter the system's default settings so that, even when provided, birth dates aren't visible unless individuals explicitly set them that way.

Researchers have started studying many similar scenarios. Privacy economics continues to evolve since its inception 40 years ago. Its combination with psychologically and behaviorally informed streams of research might prove a powerful tool to understand, and assist, privacy decision making in our complex information societies. □

References

1. T. Clark-Flory, "30 Reasons Girls Should Call It a Night," Salon.com, 5 Nov. 2007; www.salon.com/mwt/broadsheet/2007/11/05/facebook/index.html.
2. R.A. Posner, "The Economics of Privacy," *Am. Economic Rev.*, vol. 71, no. 2, 1981, pp. 405–409.
3. G.J. Stigler, "An Introduction to Privacy in Economics and Politics," *J. Legal Studies*, vol. 9, 1980, pp. 623–44.
4. H.R. Varian, "Economic Aspects

- Of Personal Privacy," *Privacy and Self-Regulation in the Information Age*, National Telecommunications and Information Administration, 1996.
5. E.M. Noam, "Privacy and Self-Regulation: Markets for Electronic Privacy," *Privacy and Self-Regulation in the Information Age*, National Telecommunications and Information Administration, 1996.
6. R.S. Laufer and M. Wolfe, "Privacy As A Concept And A Social Issue: A Multidimensional Developmental Theory," *J. of Social Issues*, vol 33, no. 3, 1977, 22–42.
7. M.J. Culnan and P.K. Armstrong, "Information Privacy Concerns, Procedural Fairness, and Impersonal Trust: An Empirical Investigation," *Organization Science*, vol. 10, no. 1, 1999, 104–115.
8. I. Simonson and A. Tversky, "Choice in Context: Tradeoff Contrast and Extremeness Aversion," *J. Marketing Research*, vol. 29, no. 3, 1992, pp. 281–295.
9. P. Slovic, "The Construction of Preference," *American Psychologist*, vol. 50, no. 5, 1995, pp. 364–371.
10. A. Acquisti, "Privacy in Electronic Commerce and the Economics of Immediate Gratification," *Proc. ACM Electronic Commerce Conf. (ACM EC 04)*, ACM Press, 2004, pp. 21–29.
11. A. Acquisti, L. John, and G. Loewenstein, "What is Privacy Worth?," tech. report, Heinz College, Carnegie Mellon University, 2009.
12. L. Brandimarte, A. Acquisti, and G. Loewenstein, "Privacy Concerns and Information Disclosure: An Illusion of Control Hypothesis," poster, *iConference 2009*; http://nora.lis.uiuc.edu/images/iConferences/Privacy_concerns_and_information_disclosure.pdf.
13. S.T. Margulis, *Psychology and Privacy in Contours of Privacy*, D. Matheson, ed., Cambridge Scholars Publishing, 2009; www.idtrail.org/files/Margulis_Contours.pdf.
14. L. Cranor, "A Framework for Reasoning about the Human in

- the Loop,” *Proc 1st Conf. Usability, Psychology, and Security*, Usenix Assoc., 2008, pp. 1–15.
15. B. Schneier, “The Psychology of Security,” *Progress in Cryptology—Proc. 1st Int’l Conf. Cryptology in Africa (AFRICACRYPT 08)*, LNCS 5023, Springer-Verlag, 2008, pp. 50–79.
 16. G. Loewenstein and E. Haisley, “The Economist as Therapist: Methodological Ramifications of ‘Light’ Paternalism,” *Perspectives on the Future of Economics: Positive and Normative Foundations*, Oxford Univ. Press, 2007.
 17. R.H. Thaler and C.R. Sunstein, *Nudge: Improving Decisions About Health, Wealth, and Happiness*, Yale Univ. Press, 2008.
 18. A. Acquisti and R. Gross, “Predicting Social Security Numbers from Public Data,” *Proc. Nat’l Academy of Science*, vol. 106, no. 27, 2009, pp. 10,975–10,980.

Alessandro Acquisti is an associate professor at the Heinz College at Carnegie Mellon University. His research interests include the economics and behavioral economics of privacy and information security, and privacy in online social networks. Acquisti has a PhD in information systems from UC Berkeley. Contact him at acquisti@andrew.cmu.edu; www.heinz.cmu.edu/~acquisti/.