PRIVACY AND MARKET FAILURES: THREE REASONS FOR CONCERN, AND THREE REASONS FOR HOPE

ALESSANDRO ACQUISTI*

INTRODUCTION

The opening panel at the University of Colorado—Boulder's "Economics of Privacy" Conference was asked to tackle an important but perilous question: Is there a market failure for information privacy? The question is perilous, because the term "market failure" is unfortunately used, and misused, to refer to different things (from market outcomes that are not Pareto efficient, correctly; to, incorrectly, any market outcome one happens not to like); but so is also (notoriously) the term "privacy."[1] The question, however, is also important: it calls us to consider whether market forces can adequately "protect" information privacy—and, in turn, what should be the essence of such protection, and what level of protection may be considered adequate. Hence, the initial query can be rephrased as: Will market forces be able to maintain a desirable balance between privacy and disclosure, in a world where most of our personal and professional lives unfold trails of electronic data, and where powerful economic interests favor information availability over information protection?

In principle, a balance between information access and information protection may be the shared long-term interest of both data subjects and data holders—more so than either extremes (unfettered access to individual data, or complete blockage of any flow of personal

* Associate Professor of Information Systems and Public Policy, Heinz College, Carnegie Mellon University. acquisti@andrew.cmu.edu. This essay is based on the keynote presented at the Silicone Flatirons "Economics of Privacy" Symposium in December 2011, and on statements prepared for the "Insights on Privacy" Roundtable at the Office of the Privacy Commissioner of Canada, February 2011.

1. *See generally* Daniel J. Solove, *A Taxonomy of Privacy*, 154 U. PA. L. REV. 477 (2006).

information). In practice, an unprecedented amount of personal information is nowadays in the hands of third parties, often out of reach, control, and even knowledge to the subjects the data refers to. That information can be used for great benefit—of both data subjects and data holders. However, vast amounts of personal data accumulated by third parties—combined with their ability to mine that data to predict behavioral patterns—can also tilt the balance of economic and social power between subjects and holders of data.

The long-term economic and social consequences of those changes in the balance of power are hard to predict. The economic theory of privacy does not provide us with conclusive answers because—with minor changes to a model's assumption—one can prove opposite conclusions with equal ease (for instance, that competitive privacy equilibria will be efficient, or, in fact, inefficient; or that privacy regulation will be redistributive, or, in fact, increase fairness).[2] The empirics of privacy cannot help us much either: as noted elsewhere, "the only straightforward conclusion about the economics of privacy and personal data is that it would be futile to attempt comparing the aggregate values of personal data and privacy protection, in search of a 'final,' definitive, and all-encompassing economic assessment of whether we need more, or less, privacy protection. Privacy means too many things, its associated trade-offs are too diverse, and consumers' valuations of personal data are too nuanced."[3]

The goal of this note is therefore much narrower. It documents some of the trends in the area of privacy that are reason for concern as they suggest rising imbalances between subjects and holders of personal data. It also documents other trends that are reason for hope as they point at ways in which the benefits of information sharing can be enjoyed while nevertheless protecting personal information.

SOME REASONS FOR CONCERN

A first reason for concern resides in the arguably unprecedented access that third parties have to aspects of individuals' lives that used to, up to not too long ago, be private. Now, those aspects are either overtly or covertly monitored by data holders, or in fact publicly broadcasted by the individuals themselves. Firms and governmental organizations have, for a long time, gathered personal information about customers and citizens. What appears remarkable, today, is the amount and quality of

2. *See* the review of competing economic theories of privacy in Section 2 of ALESSANDRO ACQUISTI, ORGANISATION FOR ECONOMIC CO-OPERATION AND DEVELOPMENT THE ECONOMICS OF PERSONAL DATA AND THE ECONOMICS OF PRIVACY 3-6 (2010), *available at* http://www.oecd.org/dataoecd/8/51/46968784.pdf.

3. *See id.* at 19.

that information, how pervasive is its collection, how invisible such collection is to the data subject, and what remarkably precise (and sometimes sensitive) inferences can be made out of that data. For instance, a few pieces of personal (but not necessarily identifiable) information can uniquely identify an individual, or allow the inference of more sensitive information about her. In a paper published in 2009, we showed how we could predict individuals Social Security numbers (in the U.S., highly sensitive information) from information gained from publicly available Internet sources.[4] As we explained in the article, we extracted birth information from Facebook profiles of students at a North American university. Then, we used simple statistical tools (such as regression analysis) to interpolate the information coming from the students' sample with information coming from the so-called Death Master File (a database of deceased individuals' Social Security numbers). Using this method, we were able to accurately predict with a single attempt the first 5 digits of the Social Security numbers for 6.3% of our sample. This result is merely one example among many of the increasing ability to predict highly sensitive data combining disparate databases, each of them not particularly sensitive. In a follow-up study, which we presented in 2011, we showed that we could combine these types of inferences together with photos from social networking sites and off-the-shelf facial recognition technology, and end up with sensitive predictions (such as someone's Social Security number) merely starting from an anonymous face.[5]

Some argue that giving users more control over their data is a way to address the above (and similar) concerns over the gathering and analysis of personal information.[6] There are reasons to be skeptical, unfortunately, that more user control, alone and by itself, can be of help. First of all, users are often unaware of the extent to which information about them is gathered and sensitive inferences are possible.[7] More importantly, while control is a *normatively* appropriate concept for privacy (that is, in terms of how we would like the world to be), the implications of control in *positive* terms (that is, in terms of how the

4. Alessandro Acquisti & Ralph Gross, *Predicting Social Security Numbers from Public Data*, 106 PROC. OF THE NAT'L ACAD. OF SCI. , 10975 (2009).

5. Alessandro Acquisti, Ralph Gross & Gred Stutzman, Presentation at BlackHat USA 2011 Conference: Faces of Facebook: Privacy in the Age of Augmented Reality (Aug. 2011), https://www.blackhat.com/docs/webcast/acquisti-face-BH-Webinar-2012-out.pdf.

6. Consider, for instance, comments by Facebook's CEO on privacy controls as instruments for increased confidence over what is shared on Facebook. Mark Zuckerberg, *Giving You More Control*, THE FACEBOOK BLOG (Oct. 6, 2010, 12:13 PM), http://www facebook.com/blog.php?post=434691727130.

7. Alessandro Acquisti & Jens Grossklags, *Privacy and Rationality in Individual Decision Making*, IEEE SEC. & PRIVACY, Jan./Feb. 2005, at 26-33.

world actually is) may be less benign. In a forthcoming manuscript,[8] we investigated how control on the publication of personal information can affect individuals' propensity to reveal sensitive details to strangers. Our conjecture was that control over publication of private information may decrease individuals' privacy concerns, and therefore increase their propensity to disclose sensitive information—even though the objective risks associated with such disclosures were more significant. To test this hypothesis, we designed a series of experiments in which we asked subjects to answer sensitive and non-sensitive questions in a survey. Across the experimental conditions, we manipulated the participants' control over information publication, but left constant (or manipulated in the opposite direction) their level of control over the actual access to and usage by others of the published information—arguably, the actual source of privacy harm. We found, paradoxically, that more control could lead to "less privacy," in the sense that higher perceived control over information publication increased our subjects' propensity to disclose sensitive information, even when the probability that strangers will access and use that information increased. These types of results show how technologies that make us feel more in control over our personal information may, in fact, promote more sensitive disclosures. These conclusions, therefore, cast some doubts over the hope that merely giving more control to users will help them achieve the desired balance between information sharing and information protection.

A third reason for concern that our recent research has highlighted relates to the impact of information about us on *others'* judgments and behaviors. In a series of experiments, we tested the hypothesis that the impact of personal information with negative valence about an individual may tend to fade away *more slowly* than the impact of information with positive valence. This would happen not just because the immediate impact of negative information may be stronger (something already shown in the literature),[9] but also because negative and positive information may be actually discounted differently.[10] In our experiments, we manipulated the type of information referring to an individual that our subjects are exposed to (namely, either positive or negative information, such as the subject engaging in a good or in a bad deed). We

8. Laura Brandimarte, Alessandro Acquisti & George Loewenstein, *Misplaced Confidences: Privacy and the Control Paradox*, J. SOC. PSYCHOL. & PERSONALITY SCI. (forthcoming), *available at* http://weis2010.econinfosec.org/papers/session2/weis2010_brandimarte.pdf.

9. *See, e.g.*, Roy F. Baumeister, Ellen Bratslavsky, Catrin Finkenauer & Kathleen D. Vohs, *Bad Is Stronger Than Good*, 5 REV. OF GEN. PSYCHOL., at 323 (2001).

10. Laura Brandimarte, Alessandro Acquisti & Joachim Vosgerau, Presentation at Workshop on the Econ. of Info. Sec.: Negative Information Looms Longer than Positive Information, (June 14, 2011), http://weis2011.econinfosec.org/program html.

also manipulated the time to which such information supposedly referred (that is, the time at which the event associated with the information ostensibly occurred: for instance, having engaged in a good/bad deed either in the recent, or in the distant, past). Then, we measured how *other* subjects reacted to such information. For instance, in one experiment we measured how the subjects judged the individual, as functions of whether the information reported about them had positive or negative valence, and whether it was presented as recent or old information. Our results confirmed that the negative effects on other people's opinion of a person, based on personal information about that person with negative valence, faded away more slowly than the positive effects of information with positive valence. In other words: good deeds positively affected our subjects' judgment of the individual only if they were reported as happening recently, and not in the past; instead, *bad* deeds negatively affected our subjects' judgment of the individual regardless of whether they had been reported as happening recently or not in the past. The implication of these results for contemporary privacy is straightforward, and rather gloomy: Web 2.0 applications allow Internet users to share all sorts of information about themselves, both positive and bad (for instance, information that may be embarrassing or inappropriate when taken out of context); the Internet not only doesn't allow that information to be "forgotten," it also seems that our innate reactions often do not allow us to "forgive" bad information about others even when it is old.

SOME REASONS FOR HOPE

Some of the cognitive and behavioral biases we investigate in the field of privacy decision making raise concerns over our ability to optimally navigate issues of privacy in the digital age. Some, however, also offer reasons for optimism.

One first reason resides in the observation that, although modern information technology seems to privilege disclosure over privacy, both the need for publicity and the need for privacy may be innate human needs—they may be part of human desires and drives across diverse times and cultures. Not only is there historical and ethnographic evidence of the quest for privacy across different societies,[11] but there is also experimental evidence suggesting that the desire to disclose and the desire to protect can be, in fact, activated through subtle manipulations. In a recent set of studies,[12] we manipulated the salience of information

11. *See* PHILOSOPHICAL DIMENSIONS OF PRIVACY: AN ANTHOLOGY (Ferdinand David Schoeman ed., 1984) (See, in particular, the chapters by Robert Murphy and Alan Westin.).

12. Leslie John, George Loewenstein & Alessandro Acquisti, Disclosure Desire: Understanding when People Disclose Private Information (unpublished manuscript on file with author).

revelation and the activation of the drive to disclose versus the drive to protect one's privacy, resulting in profoundly different effects on disclosure. Our preliminary results do suggest that individuals face competing forces when deciding how to balance information protection and disclosure (the desire to divulge, and the desire for privacy). To understand variation in information revelation across situations, we must understand how both motives operate. This, in turn, suggests that the act of disclosing plenty of personal information online does *not* prove, per se, a lack of privacy concerns.

A second reason for hope is that research on the hurdles of privacy decision-making can actually be used to develop policies and technologies that anticipate and counter those very cognitive and behavioral biases that hamper users' privacy decision-making. Such approaches are inspired by the behavioral economics literature on soft, or asymmetric, paternalism. As discussed in a recent paper,[13] research on soft paternalism suggests that lessons learned about the psychological processes underlying behavior can be used to actually aid that behavior. Systems or laws can then be designed to enhance, or even influence choice, without restricting it. The goals of these "nudging" interventions, in the privacy space, would be to increase individual and societal welfare, helping users make privacy (as well as security) decisions that they do not later regret. In doing so, this effort goes beyond privacy usability, and actually attempts to counter or anticipate the biases that lead individuals to make decisions that reduce their overall welfare or satisfaction.

An additional reason for hope resides in the development of privacy enhancing technologies (PETs). At least in principle, PETs could produce a non-zero-sum economic game between the interests of data subjects and data holders.[14] Information technologies are used to track, analyze, and link vast amounts of data about an individual, but they can also be used to aggregate, anonymize, and ultimately protect those data in ways that are both effective (in the sense that re-identifying individual information becomes too costly and therefore unprofitable) and efficient (in the sense that the desired transaction—such as an online payment, or even targeted advertising—can still be completed even though a class of individual data remains unavailable to the data holder, the merchant, or the third party). Indeed, much cryptographic research (in areas such as homomorphic encryption, secure multi-party computation, or blind signatures) could—hopefully soon—be leveraged to satisfy both needs for data sharing and needs for data privacy. Protocols to allow privacy

13. Alessandro Acquisti, *Nudging Privacy: The Behavioral Economics of Personal Information*, IEEE SEC. & PRIVACY 82 (Nov./Dec. 2009).

14. *See* ACQUISTI, *supra* note 2, at 6-7.

preserving transactions of all types (payments, browsing, communications, advertising, and so forth) have been developed. The hope is that research in this area will not stop, but in fact accelerate, so that those protocols will progress to the point where they can be cost-effectively deployed and resiliently operate in consumers' products: a future in which privacy by design and by default minimally interfere with the benefits that can be extracted from the analysis of individuals' data.

Achieving that goal will require more than self-regulation and technological ingenuity, however. It will require direct policy intervention, and will rely on our society's collective call for a future in which the balance of power between data subjects and data holders is not so dramatically skewed, as current technological and economic trends are suggesting it may be.