Editor: Alessandro Acquisti, acquisti@andrew.cmu.edu

Gone in 15 Seconds:
The Limits of Privacy Transparency and Control

Alessandro Acquisti, Idris Adjerid, and Laura Brandimarte |
Carnegie Mellon University

"I speak not to disprove what Brutus spoke, / But here I am to speak what I do know," says Antony in Shakespeare's *Julius Caesar*. We speak not to disprove transparency and control as necessary tools for consumer privacy management. But here we are to report on a rapidly amassing body of empirical evidence pointing at fundamental barriers that undermine "transparency and control" solutions to consumer privacy hurdles.

In recent years, a consensus has seemingly emerged in the US policy debate surrounding privacy. It involves an increasing reliance on self-regulatory transparency and control approaches (also called "notice and consent" or "choice and notification" regimes) to help individuals navigate increasingly complex online privacy trade-offs. The consensus is surprisingly broad: policy makers,[1] industry,[2] and advocacy organizations[3] seem to concur on the importance of granting individuals more information and more control over what happens to their data. Thus empowered, the argument goes, they'll be better able to find their personal, desired balances between disclosure and protection.

And who could disagree? In normative terms, transparency and control are clearly desirable. Information and choice are necessary for informed consent, and—with few exceptions[4,5]—usually improve decision-making. However, in positive terms (in terms of how transparency and control affect actual behavior), there are reasons to doubt the ability of either control or transparency alone to ameliorate consumers' privacy decision-making and decrease regrets associated with oversharing.[6]

Researchers have long known about problems with confusing privacy settings or complex privacy policies.[7] However, recent studies suggest that even simpler or more usable privacy controls and notices might fail to improve users' decision-making. Control might paradoxically increase riskier disclosure; transparency might be easily muted, and its effect even arbitrarily controlled.

Transparency and choice were originally part of the Organization for Economic Cooperation and Development (OECD) privacy principles (along with other principles such as purpose specification, use limitation, and accountability).[8] Disconnected from those principles, transparency and choice might reduce to necessary but not sufficient conditions for privacy protection. Worse, they might become a case of "responsibilization"—a situation in which individuals are "rendered responsible for a task which previously would have been the duty of another … or would not have been recognized as a responsibility at all."[9]

The Control Paradox

Control can reduce risk perception and increase risk-taking. This behavior is somewhat grounded in reality—an inverse correlation often exists between control and objective risk. But the perception of control is at times illusory. People might experience a great feeling of control when they determine the release of information (what to disclose and to whom), despite significant risks that might originate from that information's recipient and are out of the sender's control.

This results in the *control paradox*: individuals who perceive more control over the publication of private information pay less attention to that information's actual accessibility, and consequent use, by others. This situation occurs even though the objective risks from privacy invasions derive from how someone's data is accessed and used, not merely from how it's published. Conversely, a perception of lack of control over information publication triggers privacy concerns. It's almost as if what

Copublished by the IEEE Computer and Reliability Societies
1540-7993/13/$31.00 © 2013 IEEE

generated privacy concerns wasn't the public release of private information per se but the lack of control over that release.

Laura Brandimarte and Alessandro Acquisti, together with behavioral economist George Loewenstein, recently conducted randomized experiments investigating the control paradox.[10] The participants took a survey about their personal information and sensitive behaviors. The researchers manipulated the subjects' control over publication or accessibility of personal information in ways that either reduced or increased the objective risks associated with personal disclosure.

The results were consistent across the experiments. Participants who perceived less control over the publication or accessibility of their personal information were less willing to answer personal questions. This was true even when the probability that strangers would access (and potentially use) those answers, as well as the risks associated with that, actually decreased. Conversely, participants who perceived more control over the publication of their personal information were more willing to answer personal questions. This was true even when the risks associated with strangers accessing and using those answers (for instance, to personally identify them) actually increased.

These findings add a wrinkle to the common view of the relationship between privacy and control. Control might indeed be a way to help individuals protect their privacy while still sharing information with a selected audience. However, it might sometimes backfire, leading to riskier disclosures. Tools, interfaces, and settings ostensibly designed to protect users might create illusory control, actually exacerbating the risks they might face.

Transparency's Limits

Privacy research has suggested that hurdles in privacy decision-making might be due, at least partly, to incomplete and asymmetric information. Consumers facing privacy decisions might be unaware of how their data is collected and used, and with what consequences. Unfortunately, research has also demonstrated that the traditional tools for dealing with incomplete information—privacy policies—ineffectively communicate privacy risks. Privacy policies are often hard to find, difficult to understand, and even misinterpreted as implying, just by their mere presence, protection.[7]

Control might paradoxically increase riskier disclosure; transparency might be easily muted, and its effect even arbitrarily controlled.

To address that issue, researchers have tried to improve privacy policies' readability and usability. (For example, Patrick Kelley and his colleagues developed a nutrition-label-style presentation of policies.[11]) Simpler notices do satisfy an important informational need. But what if even such valid attempts couldn't produce "better" privacy decision-making (decisions that consumers will less likely regret or that better reflect their stated preferences)? These failures might be due to biases in how people interpret and act on available information such as privacy notices. These biases include framing effects, the use of shortcuts and heuristics for decision-making, and limited attention.[12,13]

In a series of experiments, Idris Adjerid and his colleagues investigated to what extent cognitive biases and bounded attention might limit the effectiveness of improved transparency.[14] Specifically, they examined whether even simple, accessible privacy notices might be manipulated to produce predictable, systematic changes in individuals' disclosure behavior.

In one of the experiments, students at a North American university answered questions related to various aspects of student life, some of which were potentially sensitive (for example, academic cheating or plagiarizing). All the participants received simple, short privacy notices. However, the researchers randomly assigned the participants to two groups. One group's notice said that only university students could access the answers; the other group's notice said that both students and faculty could access the answers. These notices immediately preceded the survey questions and produced the expected effect: the participants who had been told that faculty could access their answers were significantly less likely to answer the more sensitive questions. In other words, the privacy notices worked as intended.

In another experiment, the researchers introduced simple misdirections of the participants' attention for both the group told that only university students could access the answers and the group told that students and faculty could access the answers. For instance, a 15-second delay (accompanied by a timer bar) occurred between the notice's appearance and the survey questions' appearance. This delay was arguably much shorter than the delay between when Internet users read privacy policies and when they need to make privacy-sensitive decisions. However, it was enough to nullify the difference in disclosure across groups that was originally elicited by the announcement of faculty access. Other simple misdirections (such as asking participants an irrelevant question just before

they received the questionnaire) produced the same result, muting the more "sensitive" privacy notice's inhibitory effect.

These findings provide evidence of potentially systematic limitations to privacy notices' impact on disclosure. Increasing privacy policies' readability and usability can provide grounds for more informed consent. However, inconsistent decision-making might still result in the continued disparity between consumer concerns and disclosure behavior. In fact, framing or default settings might have a much more powerful role in affecting and predicting individuals' disclosure behavior than transparency or control.

So, what can be done to alleviate some of the burden and costs consumers face in managing their privacy online? Two solutions seem promising. The first is policy frameworks that reflect the original OECD privacy principles in their entirety: from purpose specification and individual participation, to use limitation and accountability. The second is interventions (including nudges[15]) aimed at anticipating and countering known hurdles and limitations that individuals face when making privacy decisions. ■

Acknowledgments

This article is based on research supported by the US National Science Foundation (NSF) under grant CNS-1012763 (Nudging Users toward Privacy), the IWT SBO (the Flemish Program for Strategic Basic Research) Security and Privacy in Online Social Networks project, and Google under a Focused Research Award on privacy nudges. This research has also been supported by NSF grants CNS-0627513 and CNS-0905562, and by Carnegie Mellon University's CyLab under US Army Research Office grants DAAD19-02-1-0389 and W911NF-09-1-0273.

References

1. *Protecting Consumer Privacy in an Era of Rapid Change*, US Federal Trade Commission, 2012; www.ftc.gov/os/2012/03/120326privacyreport.pdf.
2. R. Santalesa, "What's Next for the FTC's Proposed Privacy Framework?," Information Law Group, 2011; www.infolawgroup.com/2011/03/articles/data-privacy-law-or-regulation/whats-next-for-the-ftcs-proposed-privacy-framework.
3. R. Reitman, "FTC Final Privacy Report Draws a Map to Meaningful Privacy Protection in the Online World," Electronic Frontier Foundation, 26 Mar. 2012; www.eff.org/deeplinks/2012/03/ftc-final-privacy-report-draws-map-meaningful-privacy-protection-online-world.
4. B. Schwartz, *The Paradox of Choice: Why More Is Less*, HarperCollins, 2004.
5. J.S. Downs, G. Loewenstein, and J. Wisdom, "The Psychology of Food Consumption: Strategies for Promoting Healthier Food Choices," *Am. Economic Rev.*, vol. 99, no. 2, 2009, pp. 1–10.
6. Y. Wang et al., "I Regretted the Minute I Pressed Share: A Qualitative Study of Regrets on Facebook," *Proc. 7th Symp. Usable Privacy and Security* (Soups 11), ACM, 2011, article 10.
7. A. McDonald and L. Cranor, "The Cost of Reading Privacy Policies," *I/S: A J. of Law and Policy for the Information Society*, vol. 4, no. 3, 2009, pp. 543–568.
8. *OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data*, Org. for Economic Cooperation and Development, 23 Sept. 1980.
9. A. Wakefield and J. Fleming, *The Sage International Dictionary of Policing*, Sage Publications, 2009.
10. L. Brandimarte, A. Acquisti, and G. Loewenstein, "Misplaced Confidences: Privacy and the Control Paradox," *Social Psychological and Personality Science*, vol. 4, no. 3, 2013, pp. 340–347; http://spp.sagepub.com/content/early/2012/08/08/1948550612455931.full.pdf.
11. P.G. Kelley et al., "A 'Nutrition Label' for Privacy," *Proc. 5th Symp. Usable Privacy and Security* (Soups 09), ACM, 2009, article 4; http://cups.cs.cmu.edu/soups/2009/proceedings/a4-kelley.pdf.
12. D. Kahneman and A. Tversky, "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, vol. 47, no. 2, 1979, pp. 263–291.
13. H.A. Simon, "A Behavioral Model of Rational Choice," *Q. J. Economics*, vol. 69, no. 1, 1955, pp. 99–118.
14. I. Adjerid et al., "Sleights of Privacy: Framing, Disclosures, and the Limits of Transparency," to be published in *Proc. 9th Symp. Usable Privacy and Security* (SOUPS 13), ACM, 2013.
15. A. Acquisti, "Nudging Privacy: The Behavioral Economics of Personal Information," *IEEE Security & Privacy*, vol. 7, no. 6, 2009, pp. 82–85.

Alessandro Acquisti is an associate professor of information technology and public policy at Carnegie Mellon University. Contact him at acquisti@andrew.cmu.edu.

Idris Adjerid is an assistant professor at the University of Notre Dame Mendoza College of Business. He previously was a PhD student in privacy and technology at Carnegie Mellon University. Contact him at iadjerid@andrew.cmu.edu.

Laura Brandimarte is a postdoctoral fellow in public policy at Carnegie Mellon University. Contact her at lbrandim@andrew.cmu.edu.

cn *Selected CS articles and columns are also available for free at http://ComputingNow.computer.org.*