Social Psychological and
Personality Science
4(3) 340-347

DOI: 10.1177/1948550612455931
spps.sagepub.com

Misplaced Confidences: Privacy and the Control Paradox

Laura Brandimarte[1], Alessandro Acquisti[1] and George Loewenstein[2]

Abstract

We test the hypothesis that increasing individuals' perceived control over the release and access of private information—even information that allows them to be personally identified—will increase their willingness to disclose sensitive information. If their willingness to divulge increases sufficiently, such an increase in control can, paradoxically, end up leaving them more vulnerable. Our findings highlight how, if people respond in a sufficiently offsetting fashion, technologies designed to protect them can end up exacerbating the risks they face.

Keywords

privacy, control, paradox, behavioral economics of privacy, Web 2.0 applications

A key concern in debates about privacy is whether people are able to navigate issues of sharing and protecting personal information to their own advantage. The general assumption, which we endorse, is that policy intervention is motivated to the extent that people are poor navigators. Much as seat belts in cars are justified by the fact that people's natural driving habits (as well as those of other drivers) create an unacceptable level of risk, privacy interventions can be justified by similar limitations of individuals' abilities to manage privacy-related risks. Indeed, in recent years, considerable evidence has emerged that individuals' privacy decision making is far from optimal and is subject to various nonnormative influences. For example, privacy assurances can have the perverse effect of causing people to "clam up," whereas cues that divulgence could be risky, such as a survey's informal feel, can cause them to reveal information exactly when the situation warrants self-protective concealment of information (John, Acquisti, & Loewenstein, 2011).

The analogy to seatbelts, however, raises an important caveat, which is the central focus of the current article. Although seatbelts certainly save lives, they do not save as many lives as would be expected based on the effectiveness of the technology itself, and, research suggests, they have increased fatalities among pedestrians and bicyclists (Semmens, 1992; Wardlaw, 2000). The reason is that people who wear seatbelts tend to drive more recklessly. More generally, people often respond to safety measures intended to protect them in ways that counteract the protection—a phenomenon known as risk homeostasis or, more colloquially, the Peltzman effect (Peltzman, 1975).

In this article, we explore an analogous phenomenon in the realm of privacy. In response to the common perception that consumers are increasingly concerned about their privacy, particularly in today's Internet age, industry organizations, policy makers, and even privacy advocates have promoted solutions that involve giving individuals more control over their personal information. Consistent with a Peltzman effect, however, we document a "control paradox" such that people who experience more perceived control over limited aspects of privacy sometimes respond by revealing more information, to the point where they end up more vulnerable as a result of measures ostensibly meant to protect them. On the other hand, lower perceived control can result in lower disclosure, even if the associated risks of disclosure are lower.

Prior research has identified control as a determinant of risk perception and risk taking (e.g., Harris, 1996; Klein & Kunda, 1994; Nordgren, Van der Pligt, & Van Harreveld, 2007; Slovic, 1987; Weinstein, 1984): people are more willing to take risks, and judge those risks as less severe, when they feel in control. For example, people feel safer driving than flying, and as a result substitute road for air travel, in part based on the feeling that they have more control when driving. Such feelings are, in fact, often merited; people *do* have greater control over the risks they face in driving than they do over the risks they face

[1] H. John Heinz III College, Carnegie Mellon University, Pittsburgh, PA, USA
[2] Department of Social and Decision Sciences, Carnegie Mellon University, Pittsburgh, PA, USA

Corresponding Author:
Laura Brandimarte, H. John Heinz III College, Carnegie Mellon University, 4800 Forbes Avenue, Pittsburgh, PA 15213, USA.
Email: lbrandim@andrew.cmu.edu

in flying. However, driving is much more dangerous than flying, even for those who take exceptional measures to control their driving risks, because there are sources of risk that cannot be controlled, such as the larger number of vehicles driven and the behavior of other drivers. The ability to control *some* risks, therefore, seems to, in effect, obscure people's awareness of or attention to other risks that they cannot control.

We argue that a similar misleading feeling of control underlies many instances of problematic divulgence of information, such as the publication of embarrassing or even self-incriminating information by users of online social networks, the use of social network sites by employees to denigrate their employers, and the sharing of compromising pictures on Twitter (including the notorious case of one politician). Providing control over personal information allows one to choose how much to reveal about oneself and to whom. However, much as drivers may underestimate sources of risk that do not depend on their behavior, people who feel in control of their disclosures may underestimate the level of risk that arises from other people's access and uncontrollable usage of their disclosed information, and respond by disclosing more. On the other hand, people who feel less in control of their disclosures may overestimate those risks and respond by disclosing less.

To investigate the relationship between control, disclosure, and privacy concerns, we conducted three survey-based experiments in which respondents were asked questions that varied in sensitivity. In these experiments, we decreased (Studies 1 and 2) or increased (Study 3) participants' control over the *release* or *accessibility* of personal information. We found that perception of control affected individuals' concern about privacy to the point that their willingness to disclose sensitive information increased, even in cases where objective risks of disclosure increased. Vice versa, lack of perceived control raised privacy concerns and caused individuals to withhold information, even in cases where objective risks of disclosure decreased.

Prior empirical studies of privacy that address issues of control have shown that lower perceived control over personal information is associated with higher privacy concerns (Hoadley, Xu, Lee, & Rosson, 2010; Xu, 2007) and that individuals who are unconcerned about privacy often explain their lack of concern by noting that they feel in control of the information they reveal (Acquisti & Gross, 2006). However, to the best of our knowledge, this article is the first to demonstrate that provision of control can have a paradoxical effect: Providing users of modern information-sharing technologies with more granular privacy controls may lead them to share more sensitive information with larger, and possibly riskier, audiences.

Background and Hypotheses

In the privacy literature, "control" is construed as instrumental to privacy protection—so much so that privacy itself is often defined as the control over personal information flows (e.g., Kang, 1998; Solove, 2006; Tavani & Moor, 2001). To understand the paradox of control, however, a distinction must be drawn between the *release* of personal information (the action of willingly sharing some private information with a set of recipients), *access* to it, and *usage* by others. Disclosure—releasing personal information—is a necessary precondition for the access, use, and potential misuse of personal information by others. However, the actual costs associated with the release of personal information depend on whether other people actually access the information, and, if so, what they do with the information they access. Like the proverbial tree that falls in a forest, a violation of privacy on the Internet requires more than the posting of information: someone has to actually access and use it. For example, Facebook provides a strong feeling of control, because users can change every detail of their default privacy settings, including what type of information will be available to whom. However, users have very little control over the way in which information, once posted, will be used by a third-party application or by their friends. The third-party application could, for instance, use that information to send invasive targeted advertising to the user, or perhaps for price discrimination (Acquisti & Varian, 2005); a friend could post the information somewhere else, making it accessible to unintended third parties.

Logically, the aspect of control that should be most relevant for a decision to reveal information is control over the usage of information, since once information is released, this is the form of control that would enable the divulger to limit any negative consequences. That is, logically people should ask themselves, "If I release the information, what is likely to happen to it?" However, research on bounded rationality (Simon, 1982) and level-k thinking (Crawford, 2003; Ho, Camerer, & Weigelt, 1998; Nagel, 1995; Stahl & Wilson, 1994) shows that people often fail to engage in conditional thinking. To the degree that people fail to do so (i.e., *not* asking themselves the question of what might happen to information if they were to release it), they may focus on the most proximate level of control they have—control over release—at the expense of contemplating the actual consequences of information access and usage. Based on this logic, we predicted that people who had not yet decided whether to reveal information would fail to appreciate that control over access and usage is much more relevant than control over release once information has, in fact, been released.

Moreover, research on limited attention (e.g., Broadbent, 1957, 1982; Dukas, 2004; Johnston & Dark, 1986; Lachter, Forster, & Ruthruff, 2004; Neisser, 1967; Pashler, 1998) suggests that the human cognitive system has limited capacity and is unable to process the vast amount of information that it constantly receives. Information processing has to be selective, so when emphasis is put on a primary task, attention to secondary tasks tends to decline (Kahneman, 1973). This logic suggests that focusing people's attention on information about their level of control (or lack thereof) over release of personal data is likely to distract them from the lack of control they have over the usage that other people make of the information.

Finally, the release of information, and the choice of recipients originally intended to access it, is what people have control over, whereas actual usage involves actions by others. If people tend to overestimate the importance of their own actions

relative to others', a phenomenon documented by a large body of psychological research on "egocentrism" and perspective taking (e.g., Gilovich, Medvec, & Savitsky, 2000; Epley, Keysar, Van Boven, & Gilovich, 2004; Galinsky, 2002; Galinsky, Magee, Inesi, & Gruenfeld, 2006), they will, again, focus on their perceived control over release and access, rather than the more compelling source of risk introduced by the uncontrollable actions of others.

In three studies, we test whether these three effects can produce the kind of "control paradox" defined in the introduction. Studies 1 and 2 manipulate control over the *release* of information, and Study 3 manipulates control over *access* (but not usage) of information. The three studies show that perceived control over release or access of personal information can cause people to experience an illusory sense of security and, thus, release more information. Vice versa, lack of perceived control can generate paradoxically high privacy concerns and decrease willingness to disclose, even if the associated risks of disclosure may be lower. In addition, Study 2 tests whether, by focusing scarce attention on issues of control over release, individuals become less sensitive to other normatively relevant factors that serve as cues to objective privacy risks.

All three studies employ a paradigm that is almost ubiquitous in the experimental literature on privacy and information disclosure (e.g., see Joinson, Woodley, & Reips, 2007; Phelps, Nowak, & Ferrell, 2000, including most of the 39 studies reviewed in a meta-analysis by Weisband & Kiesler, 1996): they measure concern for privacy by people's propensity to answer personal questions in a survey (see, e.g. Frey, 1986; Singer, Hippler, & Schwarz, 1992).

Study 1

Study 1 examines the impact of decreasing control over the release of personal information on willingness to disclose, when this decrease is actually associated with lower probability of access or use of the information by others (and therefore, conditional on disclosure, lower objective benefits but also lower objective risks). Students at a North American university were recruited to participate in a survey, with the promise of snacks. Participants were invited to become members of a new campus-wide networking website that was supposedly slated to be launched at the end of the semester and populated with profiles automatically created with the information provided during the survey. The survey contained 40 questions, which varied in intrusiveness about the respondent's life in the city and on campus. Intrusiveness was measured in an initial survey of a separate sample of students from the same population. Instructions specified that none of the questions required an answer, but that all answers provided would be part of a profile that would appear on the website, visible to the university community only.

Design

The study was a between-subject design with two conditions. Participants in the certain publication condition were told that a profile would be automatically created for them, containing the information they provided, and that this profile would be published online once the website was completed. Participants in the uncertain publication condition were told that only half of the profiles created would be randomly selected to be published online. By inserting a random element in the publication process, the uncertain publication condition was intended to decrease participants' feeling of control over the public release of their survey answers, while actually reducing the probability of access by others. According to our hypotheses, the effect of decreased control would reduce willingness to disclose in the uncertain publication condition, even though objective costs or risks associated with disclosure were actually lower.

Results

Sixty-seven participants were assigned to the certain publication condition, and 65 to the uncertain publication condition (overall, 53% female; average age = 21.5, $SD = 2.85$).

Figure 1 shows the average response rate (percentage of questions answered averaged across participants) by level of intrusiveness of the questions. Across conditions, subjects were less likely to answer the more intrusive questions than the less intrusive ones, $t(130) = 11.41$, $p < .001$. Supporting our hypotheses, the main effect of control was significant, $F(1, 130) = 7.71$, $p < .001$. Moreover, as one would expect if control specifically influences concern about privacy, the two-way interaction between Condition and Question Intrusiveness was significant, $F(1, 130) = 32.43$, $p < 0.001$; participants with lower control over information release were significantly less willing to answer personal questions but especially so for more intrusive questions. The average response rate for intrusive questions was 80.8% in the certain publication condition and 61.5% in the uncertain publication condition, $t(130) = 4.16$, $p < .001$.

A lower response rate in the uncertain publication condition could also be attributable to diminished motivation to reveal information when it is less likely that the information will be publicly viewed; halving the probability of publication reduced not just the risks but also the benefits of disclosure. However, contrary to such an alternative account of the findings, intrusiveness had a negative effect on willingness to reveal. This suggests that participants were motivated to protect their sensitive information. Also, in the presence of diminished motivation, we should have observed lower response rates by subjects in the uncertain publication condition to questions that would take more effort to answer. We included in the survey open-ended questions regarding courses attended and enrollment programs to test this interpretation. A regression of aggregate word counts for the open-ended questions failed to reveal any statistically significant difference across the two conditions.

The results of Study 1, therefore, suggest that people respond to manipulations of control over release of personal information in a paradoxical way: Even though lower control implied lower objective risk of accessibility and usage of personal information by others, participants were less willing to

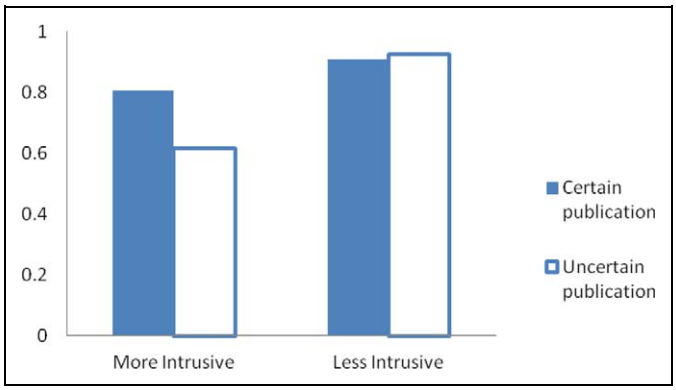

Figure 1. Average response rate by type of question in Condition 1 (filled blue, certain publication) and in Condition 2 (no-fill blue, uncertain publication)—Study 1.

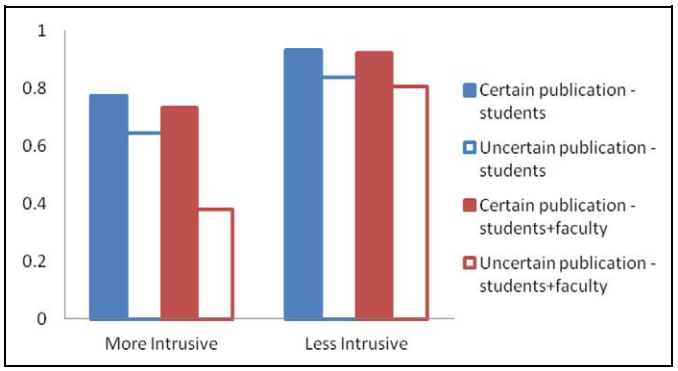

Figure 2. Average response rate by type of question in Condition 1 (filled blue; certain publication, students accessibility), Condition 2 (no-fill blue; uncertain publication, students accessibility), Condition 3 (filled red; certain publication, students + faculty accessibility) and Condition 4 (no-fill red; uncertain publication, students + faculty accessibility)—Study 2.

disclose if they were provided less control over information release.

Study 2

Design

In Study 2, we examined the impact on the propensity to answer privacy-intrusive questions of decreasing participants' perception of control over the release of personal information, while increasing the information's degree of accessibility and potential use by other, potentially more hazardous, recipients. Adopting a 2 × 2 between-subject design, Study 2 extended Study 1 by adding a between-subject manipulation of the accessibility of the information provided. University students, recruited at the same locations as Study 1, answered a shorter version of the same survey. For each of the conditions in Study 1 (certain vs. uncertain publication of participants' profiles), new conditions were created that increased accessibility by others: participants read that the website would be accessible either by students only or by students and faculty members. If one manipulation draws attention to the release of personal information, the other draws attention to its direct accessibility. The survey ended with measures of privacy and accessibility concerns, and a set of manipulation checks regarding perceived control and accessibility of the information provided.

We expected participants' willingness to disclose to be negatively affected by the accessibility of their profiles to faculty members. However, and more to the point, we expected this effect to be dampened in the case of certain publication if, consistently with the limited attention effect discussed in the introduction, participants focused on control over the release of personal information at the expense of its accessibility.

Results

Two-hundred subjects participated in Study 2 (60% female, average age = 21.3, $SD = 2.23$). Supporting our hypotheses, and replicating the results from Study 1, the main effect of control on question responding was significant, $F(1, 196) = 36.4$, $p < .001$. Moreover, similar to Study 1, there was a significant two-way interaction between Control Over Release and Question Intrusiveness, $F(1, 196) = 15.67$, $p < .001$. The main effect of accessibility by faculty was also significant, $F(1, 196) = 7.86$, $p < .01$, but, as predicted, it was smaller in the case of certain publication, as indicated by the significant interaction of control and accessibility, $F(1, 196) = 4.12$, $p < .05$. When disclosure was uncertain, participants were less willing to answer intrusive questions if the audience was composed of students and faculty as compared to students only, $t(98) = 3.92$, $p < .001$. This difference was, however, smaller and barely significant when disclosure was certain, $t(98) = 0.864$, $p = .052$. Reassurances about their control over privacy seemed to decrease participants' attention to issues of accessibility and usage (Figure 2).

Manipulation checks indicate that our manipulation of control over information release was effective, as participants perceived lower control if the publication of their profile was uncertain, $t(198) = -15.53$, $p < .001$. The main effects of control over release and accessibility were significant; when asked about accessibility concerns, participants were found to be more concerned if the publication of the profile was uncertain, $F(1, 196) = 96.51$, $p < .001$, and if the networking website was described as accessible to both students and faculty, $F(1, 196) = 15.79$, $p < .001$. The interaction was not significant. However, accessibility concerns were actually higher if the publication was uncertain and the audience was composed of students only, than in the case where the publication was certain and the audience included both students and faculty, $t(98) = 3.74$, $p < .001$.

Similar results were found for privacy concerns; participants reported higher concerns if the publication of the profile was uncertain, $F(1, 196) = 215.36$, $p < .001$, and if the networking website was accessible to both students and faculty, $F(1, 196) = 5.01$, $p < .05$. The interaction was not significant; however, reported privacy concerns were higher if the publication was uncertain and the audience was composed of students only,

than in the case where the publication was certain and the audience included both students and faculty, $t(98) = 8.50$, $p < .001$. This may suggest that privacy concerns mediate the effect of control on willingness to disclose—a conjecture that we test in Study 3.

Overall, Study 2 supports the central idea that privacy concerns are affected by control over release of personal information and that reassurances about control over release can distract people from concerns about potentially more hazardous accessibility.[1]

Study 3

In contrast to Studies 1 and 2, Study 3 tested the impact of providing participants with more, rather than less, control over the release of their information, and with more control over the actual accessibility of disclosed information. Study 3 also extended the previous studies by testing whether the effect of control applies even to the disclosure of information that could be used to personally identify the divulger, which would significantly heighten the objective risk of privacy violations (Sweeney, 1997). Finally, in Study 3 we tested whether privacy concerns mediated the effect of the experimental manipulations on willingness to disclose.

Using similar recruitment methods as the previous studies, participants were invited to take a survey on ''ethical behaviors.'' The survey consisted of 10 yes/no questions regarding more or less sensitive behaviors, such as stealing, lying, and consuming drugs. Perceived intrusiveness of the questions was established following the same procedure used in the previous studies.

Participants were informed that none of the questions required an answer and that the researchers intended to publish the results of the study—including participants' anonymous survey answers—in a research bulletin. No detail was given as to whom this bulletin would be accessible, which was a constant feature across all conditions.

Design

The study was a nonfactorial between-subject design with four conditions, characterized by increasing control over release of personal information.

In the implicit control condition, participants read that by answering a question they would automatically give the researchers permission to publish the answer provided in a research bulletin. Participants could decide *not* to answer any question and therefore deny the researchers the ability to publish their answers, but, unlike the other conditions, there was no explicit mention of the existence of such control.

In the explicit aggregate control condition, before answering the 10 questions on ethical behaviors, participants were asked to check a box if they agreed to give the researchers permission to publish *all* their answers among the results of the study. The default option was that the answers would not be published.

In the explicit granular control condition, for each individual question participants were asked to check a box, next to the question, to signal that they were willing to grant publication permission of their answer to that specific question. The default option was that the answers would not be published. This condition emulates several Web 2.0 services, such as blogs and online social networks, which provide users with granular control on what to publish online.[2]

Finally, the explicit aggregate control with demographic condition was identical to the explicit aggregate control condition but asked for permission to publish demographic information (in all of the other conditions, participants read that the demographic information they provided would not be published). Participants could click on separate publication permission boxes for gender, age, and country of birth. Releasing this type of personal information is objectively riskier than releasing only answers to ethical behaviors, as it greatly increases the risk that participants could be identified.

We expected to see larger willingness to disclose as the granularity of privacy controls increased, especially for more intrusive questions. Therefore, we predicted that willingness to disclose would be lowest in the implicit control condition and highest in the explicit granular control condition, with the remaining conditions in between. Consistent with a control paradox, we predicted that privacy concerns would be soothed by the existence of *explicit control* over access, leading to greater public disclosure of personal information.

Similar to Study 2, the survey ended with a measure of privacy and accessibility concern and a set of manipulation checks regarding perceived control.

Note that for participants with implicit control, answering a question implied the publication of the corresponding answer, while participants in all other conditions could decide to provide an answer but, when explicitly asked, grant no publication permission. Given this setup, to meaningfully compare results across all conditions, we compared the level of positive responses in the control condition to responses that participants not only provided but also consented to be published.[3]

Results

A total of 134 subjects participated in Study 3 (50% female, average age = 21.9, $SD = 2.72$).

All participants in the explicit aggregate control conditions, with and without demographics, checked the publication permission box, thus allowing the public release of their answers. Moreover, all participants in the demographic condition granted permission to publish all three demographic items—which dramatically increases their identifiability (Sweeney, 1997). This striking result suggests that, as long as people perceive control over the decision to publish personal information and the audience to whom access will be granted, they will indeed decide to publish it, even if the objective risks associated with disclosure increase dramatically. The main effect of control over information release was significant, $F(3, 130) = 33.53$, $p < .001$; Figure 3 shows that willingness to disclose

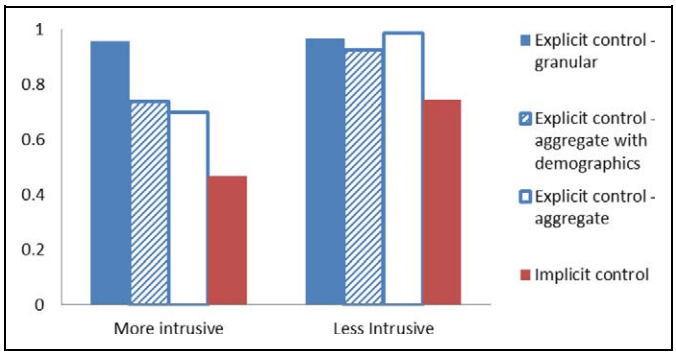

Figure 3. Average response rate by type of question in Condition 1 (filled red; implicit control on publication), in Condition 2 (no-fill blue; explicit control, aggregate), in Condition 3 (filled blue; explicit control, granular), and in Condition 4 (striped blue; explicit control, aggregate with demographics)—Study 3.

increases as the level of control increases from implicit to explicit aggregate and to explicit granular. In addition, consistent with the idea that control influences concern about privacy, the two-way interaction between Condition and Question Intrusiveness was significant, $F(3, 130) = 11.98$, $p < .001$. Supporting our hypothesis that perceived control decreases people's sensitivity to privacy violations, voluntarily revealing demographic information in the demographic condition did not affect willingness to answer sensitive questions, even though the objective risk of disclosure was higher.

These results suggest that reported privacy concerns should mediate the effect of actual control (dummy variables representing all conditions with explicit control) on willingness to disclose. To test this, we included our measure of privacy concern in a mediation analysis (Table 1 in the Online Appendix, see Online Supplemental Material found at http://spps.sagepub.com/supplemental). We conducted an ordinary least squares (OLS) regression using a bootstrapping technique (Preacher & Hayes, 2008). The total effect of actual control on willingness to disclose was positive and significant, as the coefficients on all three dummies were significantly larger than zero (Model 1: $\beta_2 = .25$, $SE = .04$, $t(130) = 6.47$, $p < .001$; $\beta_3 = .36$, $SE = .04$, $t(130) = 9.84$, $p < .001$; $\beta_4 = .24$, $SE = .04$, $t(130) = 6.25$, $p < .001$). Privacy concern correlated negatively with actual control (Model 2: $\beta_2 = -.82$, $SE = .43$, $t(130) = -1.91$, $p = .06$; $\beta_3 = -2.11$, $SE = .42$, $t(130) = -5.06$, $p < .001$; $\beta_4 = -1.82$, $SE = .43$, $t(130) = -4.26$, $p < .001$). Accounting for privacy concerns, the relationship between actual control and willingness to disclose weakened (Model 3: $\beta_2 = .22$, $SE = .04$, $t(129) = 6.08$, $p < .001$; $\beta_3 = .30$, SE = .04, $t(129) = 7.87$, $p < .001$; $\beta_4 = .18$, $SE = .04$, $t(129) = 4.78$, $p < .001$). A bootstrap analysis revealed that the 95% bias-corrected confidence interval for the size of the indirect effects excluded zero, which suggested a significant indirect effect (MacKinnon, Fairchild, & Fritz, 2007; Preacher & Hayes, 2004).

This study shows that, paradoxically, participants were more likely to allow the publication of information about them and more likely to disclose *more* information of a sensitive nature, as long as they were *explicitly*, instead of *implicitly*, given control over its publication. Participants in the implicit control condition could avoid publication by not answering questions; but participants in the other conditions, who had an explicit option to publish their answers and determine the level of their accessibility, felt less privacy concerned and thus became more likely to not just *answer* but also allow the *publication* of their answers. It was not the publication of personal information per se that modulated privacy concerns but rather the explicit perceived control over it.

Discussion

Three experiments provide empirical evidence that perceived control over release plays a critical role in sharing/oversharing personal information, relative to the objective risks associated with information access and usage by others. In Study 1, participants responded to manipulations that decreased control over information release, even though risks associated with information access and use by others were in fact decreased. In Study 2, control over release distracted participants from concerns about potentially more hazardous accessibility. In Study 3, participants given explicit control over the release and accessibility of their personal information revealed more, even exposing themselves to higher risks of identifiability.

Our findings introduce a novel scenario in the scholarly literature on privacy and control, where it has been conventional wisdom that control over personal information either implies (Culnan, 1993; Elgesem, 1996; Fried, 1984; Lessig, 2002; Miller, 1971; Smith, Milberg, & Burke, 1996; Westin, 1967) or at most does not negatively affect (Laufer & Wolfe, 1977; Tavani & Moor, 2001) privacy protection. Our results show that "more" control can sometimes lead to "less" privacy in the sense of higher objective risks associated with the disclosure of personal information. In other words, our results provide evidence that control over personal information may be a necessary (in ethical or normative terms) but not sufficient condition for privacy protection.

Notice that our argument does not posit that people *should* be concerned about their privacy, or that they *have to* disclose less in order to achieve higher utility or satisfaction. While recent research on regrets associated with online information sharing does indicate that, at times, people feel they revealed too much (Wang et al., 2011), in our studies, and indeed most situations, there is no objective standard for determining whether participants revealed too little or too much. To document that privacy-related behavior is suboptimal, therefore, we show that people change their propensity to disclose in response to nonnormative factors (such as whether they have explicit or implicit control over publication) and fail to change their disclosure behavior (or even change in the wrong direction) in response to normative factors (such as whether they can be personally identified).

The conventional wisdom that control is an essential component of privacy is so ubiquitous that control has become a code word employed both by legislators and government

bodies in proposals for enhanced privacy protection and by data holders and service providers to deflect criticisms regarding the privacy risks borne by data subjects. For instance, Facebook's CEO Mark Zuckerberg has repeatedly stressed the role of privacy controls as instruments to have "more confidence as you share things on Facebook,"[4] while both Senator Kerry's bill proposal and the recent Federal Trade Commission's Privacy Report focus on giving users more (privacy) control.[5] In fact, numerous government and corporate entities in the United States have advocated self-regulatory "choice and consent" models of privacy protection that, essentially, rely on users' awareness and control.

The argument is appealing; users *do* want more control over how their information is collected and used (Consumer Reports National Research Center, September 2008, http://www.consumersunion.org/pub/core_telecom_and_utilities/006189.html). However, higher levels of control may not always serve the ultimate goal of enhancing privacy protection. The paradoxical policy implication of these findings is that the feeling of security conveyed by the provision of fine-grained privacy controls may lower concerns regarding the actual accessibility and usability of information, driving those provided with such protections to reveal more sensitive information to a larger audience.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interests with respect to the research, authorship, and/or publication of this article.

Funding

The authors disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: IWT SBO Project on Security and Privacy for Online Social Networks (SPION), U.S. Army Research Office under Contract DAAD190210389 through Carnegie Mellon CyLab, and TRUST (Team for Research in Ubiquitous Secure Technology), which receives support from the National Science Foundation (NSF award number CCF-0424422) and the following organizations: AFOSR (#FA9550-06-1-0244), BT, Cisco, DoCoMo USA Labs, EADS, ESCHER, HP, IBM, iCAST, Intel, Microsoft, ORNL, Pirelli, Qualcomm, Sun, Symantec, TCS, Telecom Italia, and United Technologies. Acquisti also gratefully acknowledges support from Google's Focused Research Award and from National Science Foundation (Award CNS-1012763).

Notes

1. Consistent results were obtained from a similar study, in which accessibility was manipulated telling students that their online profile would be accessible to either members of their own university only or members of both their own university and of another, larger university in the same neighborhood. The results are available from the authors.
2. The original study included one additional condition, similar to Condition 3, but with the default consisting of the answers being granted publication. The purpose of that condition was to make sure that default effects were not the main driver for allowing (or not) the publication of the answers. The results are consistent with those presented in the article and are available from the authors.
3. The results obtained considering response rates (the dependent variable used for the previous studies) are similar to those obtained for publication rates and are available from the authors.
4. "Giving you more control," posted by Mark Zuckerberg on October 10, 2010, retrieved from http://www.facebook.com/blog.php?post=434691727130
5. See http://kerry.senate.gov/press/release/?id=223b8aac-0364-4824-abad-274600dffe1c and http://www.ftc.gov/os/2010/12/101201privacyreport.pdf

References

Acquisti, A., & Gross, R. (2006). Imagined communities: awareness, information sharing, and privacy on the Facebook. *Proceedings of the Sixth Workshop on Privacy Enhancing Technologies*.

Acquisti, A., & Varian, H. (2005). Conditioning prices on purchase history. *Marketing Science*, *24*, 367–381.

Broadbent, D. E. (1957). A mechanical model for human attention and immediate memory. *Psychological Review*, *64*, 205–215.

Broadbent, D. E. (1982). Task combination and selective intake of information. *Acta Psychologica*, *50*, 253–290.

Crawford, V. P. (2003). Lying for strategic advantage: Rational and boundedly rational misrepresentation of intentions. *American Economic Review*, *93*, 133–149.

Culnan, M. J. (1993). How did they get my name? An exploratory investigation of consumer attitudes toward secondary information use. *MIS Quarterly*, *17*, 341–363.

Dukas, R. (2004). Causes and consequences of limited attention. *Brain, Behavior and Evolution*, *63*, 197–210.

Elgesem, D. (1996). Privacy, respect for persons, and risk. In C. Ess (Ed.), *Philosophical perspectives on computer-mediated communication*. New York: State University of New York Press.

Epley, N., Keysar, B., Van Boven, L., & Gilovich, T. (2004). Perspective taking as egocentric adjustment. *Journal of Personality and Social Psychology*, *87*, 327–339.

Frey, J. H. (1986). An experiment with a confidentiality reminder in a telephone survey. *Public Opinion Quarterly*, *50*, 267–269.

Fried, C. (1984). Privacy. In F.D. Schoeman (Ed.), *Philosophical dimensions of privacy*. New York, NY: Cambridge University Press.

Galinsky, A. D. (2002). Creating and reducing intergroup conflict: The role of perspective-taking in affecting out-group evaluations. In H. Sondak (Ed.), *Research on managing groups and teams: toward phenomenology of groups and group membership* (Vol. 4, pp. 85–113). Greenwich, CT: Elsevier/JAI.

Galinsky, A. D., Magee, J. C., Inesi, M. E., & Gruenfeld, D. H. (2006). Power and perspectives not taken. *Psychological Science*, *17*, 1068–1074.

Gilovich, T., Medvec, V. H., & Savitsky, K. (2000). The spotlight effect in social judgment: An egocentric bias in estimates of the salience of one's own actions and appearance. *Journal of Personality and Social Psychology*, *78*, 211–222.

Harris, P. (1996). Sufficient grounds for optimism? The relationship between perceived controllability and optimistic bias. *Journal of Social and Clinical Psychology*, *15*, 9–52.

Ho, T.-H., Camerer, C., & Weigelt, K. (1998). Iterated dominance and iterated best response in experimental p-beauty contests. *American Economic Review*, *88*, 947–969.

Hoadley, C. M., Xu, H., Lee, J. J., & Rosson, M. B. (2010). Privacy as information access and illusory control: The case of the Facebook News Feed privacy outcry. *Electronic Commerce Research and Applications*, *9*, 50–60.

John, L., Acquisti, A., & Loewenstein, G. (2011). The best of strangers: Context dependent willingness to divulge sensitive information. *Journal of consumer research*, *37*(5), 858–873.

Johnston, W. A., & Dark, V. (1986). Selective attention. *Annual Review of Psychology*, *37*, 43–75.

Joinson, A. N., Woodley, A., & Reips, U. D. (2007). Personalization, authentication and self-disclosure in self-administered Internet surveys. *Computers in Human Behavior*, *23*, 275–285.

Kahneman, D. (1973). *Attention and effort*. New York, NY: Prentice Hall.

Kang, J. (1998). Information privacy in cyberspace transactions. *Stanford Law Review*, *50*, 1193–1294.

Klein, W. M., & Kunda, Z. (1994). Exaggerated self-assessments and the preference for controllable risks. *Organizational Behavior and Human Decision Processes*, *59*, 410–427.

Lachter, J., Forster, K. I., & Ruthruff, E. (2004). Forty-five years after Broadbent (1958): Still no identification without attention. *Psychological Review*, *111*, 880–913.

Laufer, R. S., & Wolfe, R. (1977). Privacy as a concept and a social issue: A multidimensional developmental theory. *Journal of Social Issues*, *33*, 22–41.

Lessig, L. (2002). Privacy as property. *Social Research*, *69*, 247–269.

MacKinnon, D. P., Fairchild, A. J., & Fritz, M. S. (2007). Mediation analysis. *Annual Review of Psychology*, *58*, 593–614.

Miller, A. R. (1971). *The assault on privacy*. Ann Arbor: The University of Michigan.

Nagel, R. (1995). Unraveling in guessing games: An experimental study. *American Economic Review 85*, 1313–1326.

Neisser, U. (1967). *Cognitive psychology*. New York, NY: Appleton.

Nordgren, L. F., Van der Pligt, J., & Van Harreveld, F. (2007). Unpacking perceived control in risk perception: The mediating role of anticipated regret. *Journal of Behavioral Decision Making*, *20*, 533–544.

Pashler, H. E. (1998). *The psychology of attention*. Cambridge, MA: MIT Press.

Peltzman, S. (1975). The effects of automobile safety regulation. *Journal of Political Economy*, *83*, 677–726.

Phelps, J., Nowak, G., & Ferrell, E. (2000). Privacy concerns and consumer willingness to provide personal information. *Journal of Public Policy & Marketing*, *19*, 27–41.

Preacher, K. J., & Hayes, A. F. (2004). SPSS and SAS procedures for estimating indirect effects in simple mediation models. *Behavior Research Methods, Instruments, and Computers*, *36*, 717–731.

Preacher, K. J., & Hayes, A. F. (2008). Asymptotic and resampling strategies for assessing and comparing indirect effects in multiple mediator models. *Behavior Research Methods*, *40*, 879–891.

Semmens, J. (1992). Do seat belt laws work? *The Freeman*, *42*. Retrieved from http://www.thefreemanonline.org/columns/do-seat-belt-laws-work/

Simon, H. A. (1982). *Models of bounded rationality*. Cambridge, MA: MIT Press.

Singer, E., Hippler, H. J., & Schwarz, N. (1992). Confidentiality assurances in surveys: Reassurance or threat? *International Journal of Public Opinion Research*, *4*, 256–268.

Slovic, P. (1987). Perception of risk. *Science*, *236*, 280–285.

Smith, H. J., Milberg, S. J., & Burke, S. J. (1996). Information privacy: Measuring individuals' concerns about organizational Practices. *MIS Quarterly*, *20*, 167–196.

Solove, D. J. (2006). *A taxonomy of privacy. University of Pennsylvania Law Review*, *154*, 477–560.

Stahl, D., & Wilson, P. (1994). Experimental evidence on players' models of other players. *Journal of Economic Behavior and Organization*, *25*, 309–327.

Sweeney, L. (1997). Weaving technology and policy together to maintain confidentiality. *Journal of Law, Medicine & Ethics*, *25*, 98–110.

Tavani, H. T., & Moor, J. H. (2001). Privacy protection, control of information, and privacy-enhancing technologies. *Computers and Society*, *31*(1), 6–11.

Wang, Y., Komanduri, S., Leon, P. G., Norcie, G., Acquisti, A., & Cranor, L. F. (2011). I regretted the minute I pressed share: A qualitative study of regrets on Facebook. *7th Symposium on Usable Privacy and Security, Pittsburgh, PA, USA*.

Wardlaw, M. J. (2000). Three lessons for a better cycling future. *British Medical Journal*, *321*, 1582–1585.

Weinstein, N. D. (1984). Why it won't happen to me: Perceptions of risk factors and susceptibility. *Health Psychology*, *14*, 431–457.

Weisband, S., & Kiesler, S. (1996). Self-disclosure on computer forms: Meta-analysis and implications. *Proceedings of the SIGCHI Conference on Human factors in computing systems, Vancouver, Canada*.

Westin, A. R. (1967). *Privacy and freedom*. New York, NY: Atheneum.

Xu, H. (2007). The effects of self-construal and perceived control on privacy concerns. *Proceedings of 28th Annual International Conference on Information Systems*, Montréal, Canada.

Author Biographies

Laura Brandimarte is a PhD candidate in public policy and management at the H. John Heinz III College, Carnegie Mellon University.

Alessandro Acquisti is an associate professor of information technology and public policy at the H. John Heinz III College, Carnegie Mellon University.

George Loewenstein is the Herbert A. Simon Professor of economics and psychology at the Department of Social and Decision Sciences, Carnegie Mellon University.