

Privacy, Trust, and Self-Disclosure Online

Adam N. Joinson,¹ Ulf-Dietrich Reips,² Tom Buchanan,³
and Carina B. Paine Schofield⁴

¹*University of Bath, UK*

²*University of Deusto and Ikerbasque, Basque Foundation of Science, Spain*

³*University of Westminster, UK*

⁴*Asbridge Business School, UK*

Despite increased concern about the privacy threat posed by new technology and the Internet, there is relatively little evidence that people's privacy concerns translate to privacy-enhancing behaviors while online. In Study 1, measures of privacy concern are collected, followed 6 weeks later by a request for intrusive personal information alongside measures of trust in the requestor and perceived privacy related to the specific request ($n = 759$). Participants' dispositional privacy concerns, as well as their level of trust in the requestor and perceived privacy during the interaction, predicted whether they acceded to the request for personal information, although the impact of perceived privacy was mediated by trust. In Study 2, privacy and trust were experimentally manipulated and disclosure measured ($n = 180$). The results indicated that privacy and trust at a situational level interact such that high trust compensates for low privacy, and vice versa. Implications for understanding the links between privacy attitudes, trust, design, and actual behavior are discussed.

1. PRIVACY, TRUST, AND THE INTERNET

The use of new technology, and particularly the Internet, increasingly requires people to disclose personal information online for various reasons. In com-

Adam Joinson is a psychologist with interests in information systems, computer-mediated communication and privacy; he is Reader in Information Systems in the School of Management at the University of Bath. **Ulf-Dietrich Reips** is a psychologist with interests in Internet-based research methodology (particularly experimental methodology); he is an Ikerbasque Research Professor at the Departamento de Psicología of the Universidad de Deusto in Bilbao. **Tom Buchanan** is a psychologist with interests in Internet behavior; he is Reader in Psychology in the Department of Psychology, University of Westminster. **Carina Paine Schofield** is a psychologist with interests in technology mediated learning; she is a researcher at the Ashridge Business School.

CONTENTS

1. PRIVACY, TRUST, AND THE INTERNET
2. STUDY 1: MODELING PRIVACY, TRUST, AND DISCLOSURE
 - 2.1. Method
 - Participants
 - Materials
 - Procedure
 - 2.2. Results
 - Descriptive Statistics
 - Correlations Between Measures
 - Privacy Attitudes, Reported Behaviours, and Self-Disclosure
 - Model Testing
 - Situational Factors and Nondisclosure
 - Situational and Dispositional Privacy and Trust and Nondisclosure
 - 2.3. Discussion
3. STUDY 2: EXPERIMENTAL MANIPULATION OF PRIVACY AND TRUST
 - 3.1. Method
 - Participants
 - Experimental Manipulations
 - Measurement of Self-Disclosure
 - Perceived Trust and Privacy
 - Procedure
 - 3.2. Results: Self-Reported Trust, Perceived Privacy, and Experimental Condition
 - 3.3. Discussion
4. GENERAL DISCUSSION
5. LIMITATIONS AND FUTURE WORK
6. CONCLUSIONS
- APPENDIX A1. HIGH PRIVACY FRONT PAGE
- APPENDIX A2. LOW PRIVACY FRONT PAGE STATEMENT

puter-mediated communication, disclosure may serve to reduce uncertainty in an interaction (Tidwell & Walther, 2002) or to establish legitimacy when joining an online group (Galegher, Sproull, & Kiesler, 1998). Disclosure is often a prerequisite to access services (via, e.g., the ubiquitous registration form) or to make online purchases (Metzger, 2006) or is requested for those same services to be personalized (e.g., in the form of recommendations or “one-click” purchasing). The increasingly social nature of much Web-based software (e.g., social network sites) also places a privacy cost on users due to an heightened requirement for disclosure of personal information as part of the functionality of the system itself (Joinson, 2008). For instance, the capacity to upload location-aware photographs from camera phones requires users to make a series of privacy-related judgements about levels of public access related to both security considerations and the risk of self and other disclosure (Ahern et al., 2007). In addition to this increased need for disclosure, the development of ambient and ubiquitous technologies has raised the possibility that devices will communicate, or even broadcast, personal information

without recourse to the user themselves (Bellotti & Sellen, 1993); the ability to easily store information, and cross-reference databases, raises the possibility of unwitting disclosure through information accrual.

Perhaps not surprisingly, this has raised a number of privacy concerns among consumers and privacy advocates alike (e.g., Ackerman, Cranor, & Reagle, 1999; Jupiter Research, 2002a; UK Information Commissioner, 2006). According to a Harris and Associates Inc. (2004) survey, 65% of respondents reported that they had declined to register at an e-commerce site because of privacy concerns. A recent poll by UPI-Zogby (2007) found that 85% of respondents said that the privacy of their personal information was important to them as consumers. The Oxford Internet Institute survey (Dutton & Helsper, 2007) found that 70% of U.K. Internet users agreed or strongly agreed with the statement that “people who go on the Internet put their privacy at risk,” and 84% (up from 66% in 2005) agreed that “personal information is being kept somewhere without my knowledge.”

Forrester Research (2005) reported that, despite evidence that consumer privacy attitudes are “all bark and no bite,” “companies that advocate for customer privacy will be more successful than those that ignore consumer concerns.” Harris (2004) also noted that although the costs of data protection regulation to business can be substantial, compliance also brings significant economic benefits. Acquisti, Friedman, and Telang (2006) reported that data breaches can impact negatively (albeit temporarily) on the stock market valuation of an organization.

However, the relationship between consumers’ privacy concerns and actual behavior neither is straightforward nor has any link been established incontrovertibly. There is evidence that although many Internet users express privacy-protectionist attitudes, this rarely translates to their actual behavior (Forrester Research, 2005, Jupiter Research, 2002b; Metzger, 2006; Pew Internet and American Life Project, 2000). For instance, Spiekermann, Grossklags, and Berendt (2001) measured the privacy preferences of 171 users and observed their behavior on a mock e-commerce site. On this site, the users were “helped” by a ’bot (short for an automated agent or “robot”) that asked a number of purchase-related questions of differing levels of intrusiveness. They found very little evidence that privacy preferences were related to people’s actual behavior in response to the ’bot’s questions. Similarly, Metzger (2006) found no association between people’s privacy concerns and their disclosure to an e-commerce site, or between the content of a privacy policy or presence of a privacy seal and disclosure behavior. The failure of various privacy enhancing technologies in the marketplace also suggests a disjunction between people’s stated attitudes and their actual actions to protect their privacy (Acquisti & Grossklags, 2003).

This lack of evidence linking privacy concerns or reassurances and actual behavior can be explained in a number of ways. First, there is evidence that people do not read privacy policies, or if they do read them do not fully understand their contents (Berendt, Gunther, & Spiekermann, 2005; Milne & Culnan, 2004). In part, this may be due to the lack of usability of the policies themselves, with the majority requiring greater than high school education to be comprehended (Jensen & Potts, 2004). There

is, however, some evidence that people may be willing to pay a premium for privacy protection when privacy information is made readily available (Gideon, Cranor, Egelman, & Acquisti, 2006; Tsai, Egelman, Cranor, & Acquisti, 2007, but see Acquisti & Grossklags, 2007), supporting the argument of Ackerman and Cranor (1999) for semiautonomous “privacy critics” to support users’ decision making, and toward efforts to establish easy to access (and understand) privacy policies (e.g., P3P; Cranor, Langheinrich, Marchiori, Presler-Marshall, & Reagle, 2002)

A second issue is that many of the studies of privacy concern and behavior have measured reported disclosure or intended disclosure rather than actual behavior (e.g. Chellappa & Sin, 2005, Malhotra, Kim, & Agarwal, 2004), although when actual behavior is measured, the same pattern of results often emerges (e.g., Metzger, 2006; Spiekermann et al., 2001, but see also Gideon et al., 2006, for contrary results).

Finally, the measurement of privacy concern tends to be generic and does not usually include previous disclosure behavior (see Metzger, 2006), nor does it often reflect the distinction between generic privacy concerns and people’s interpretation of a specific situation (Margulis, 2003, Palen & Dourish, 2003). For instance, Olson, Grudin, and Horvitz (2005) noted that “people’s willingness to share depends on who they are sharing the information with” (p. 1987). This highlights the importance of the *relationship* between the discloser and the recipient in determining disclosure behavior, including the trust we have in the other party to our information.

Trust is critical in understanding when we choose to share personal information with others and when we choose secrecy. Altman (1975) described a *self boundary* (the boundary around the person) that is modified by self-disclosure; and a *dyadic boundary* that ensures the discloser’s safety from leakage of information to uninvited third parties. The self boundary may be open or closed depending on such interpersonal factors as the level of trust in a disclosure target (Altman & Taylor, 1973). Within e-commerce, trust has been identified as a key factor in determining purchasing behavior: “If the web site does not lead the consumer to believe that the merchant is trustworthy, no purchase decision will result” (Ang & Lee, 2000, p. 3). Metzger (2006) found that trust (conceptualized as reputation) predicted disclosure behavior to a mock music CD e-commerce site.

Online privacy is often framed as a contributor to trust, rather than as an independent effect on online behavior. For instance, Google Inc. privacy counsel for Europe justified the anonymizing of search data by stating, “We believe that privacy is one of the cornerstones of trust” (Fildes, 2006). This relationship is borne out in a series of research findings. For instance, Malhotra et al. (2004) examined the links between people’s Internet information privacy concerns and their related behavioral intentions. They found that the effect of privacy concerns on behavioral intentions was mediated by trust. Similarly, Chellappa and Sin (2005) studied consumers’ intent to use personalization services. They also found that this intent was influenced by both trust and concern for privacy. Metzger (2004) also found that the link between privacy concern and self-disclosure was mediated by trust. These findings would go some way toward explaining why privacy attitudes rarely predict actual behavior because any explanatory

power of privacy concern would be mediated by trust, and because much of the research is conducted in lab-based studies, we would expect trust to be consistently high.

In its traditional sense, mediation refers to the effect of an independent variable on a dependent variable being explained by common links to a third variable (i.e., the mediator; Baron & Kenny, 1986). For instance, a correlation between income and cancer might be explained by a correlation between income and smoking (the mediator), and then between smoking and cancer. Thus, the results just reported would suggest that privacy has little or no direct effect on behaviour; instead, any effect can be explained by the links between privacy and trust, and then between trust and behavior.

However, this interpretation of the nature of the relationship between privacy and trust is potentially problematic. First, as argued previously, privacy is more dynamic than simply people's general attitudes and needs to take into account the specific interaction. Second, in interpersonal research (e.g., computer-mediated communication), there is considerable evidence that anonymity (one form of privacy) can substantially increase self-disclosure, in part because issues of trust are rendered irrelevant (e.g., Ben Ze'ev, 2003; Joinson, 2001; Rubin, 1975). If this is the case, the relationship between trust and privacy on behavior may take the form of an interaction (also called moderation). Moderation is when the impact of a variable on an outcome measure (e.g., vitamin supplements on health) is altered by the status of a moderator (e.g., vitamin plus precondition leads to negative outcome, absence of precondition leads to positive outcomes). In the case of privacy and trust, it may be that the effect of privacy on behavior is moderated by trust, such that in conditions of low trust privacy exerts an influence on behavior, whereas in high trust environments, privacy has a negligible impact on behavior.

In the present article we examine the nature of the relationship between privacy (dispositional concerns and situational), trust, and a privacy-related behavior: self-disclosure. Uniquely, self-disclosure is behaviorally measured outside of an e-commerce environment, and we utilize both survey-related and experimental methods to test for mediation and moderation. In Study 1, measures of dispositional privacy, perceived privacy, and trust are studied in light of disclosure behavior. In Study 2, privacy and trust are experimentally manipulated to test for a moderation effect, and the impact on both disclosure behavior and perceived privacy and trust examined.

2. STUDY 1: MODELING PRIVACY, TRUST, AND DISCLOSURE

In Study 1, participants completed two sets of measures across two periods. At Time 1, privacy concern (attitude and previous behavior) measures were completed. Six weeks later (Time 2), self-disclosure measures and situational perceived privacy and trust were collected. The analyses examine the links between both dispositional and situational measures and disclosure.

2.1. Method

Participants

Participants were 759 members of an online research panel of Open University (OU) students called PRESTO. The OU is an adult distance learning institution with nearly all students studying part time from home or work. PRESTO members are recruited annually and commit to completing six online surveys over 12 months. The sample is selected using stratified sampling (e.g., by age, gender, academic discipline, and geographic location). Of the 759 respondents, 64% (487) were female and 36% (272) were male. The mean age of the sample was 42.58 years, (range = 17–84 years, $SD = 11.11$).

Materials

Time 1: Privacy Dispositions. A set of 16 privacy attitude items and 12 reported privacy behavior items developed by Buchanan, Paine, Joinson, and Reips (2007) was given to participants. For all privacy items, responses were made on a 5-point scale (from *very concerned* to *not at all concerned*). The privacy behavior items consisted of 6 “general caution” items (e.g., reading privacy policies, license agreements, etc.) and 6 “technical protection” items (e.g., removing cookies, clearing Internet browser history regularly, etc.; both ranging from *always* to *never*). Participants were also asked about their Internet use (history, breadth of use, and time spent online). These data are not analyzed in the present study.

Time 2: Self-Disclosure and Situational Aspects of Privacy. Participants completed a 10-item measure of behavioral self-disclosure. In this measure, participants respond to a sensitive item such as “How many different sexual partners have you had?” using one of three options: they could submit the default option “please choose” (termed “passive nondisclosure”), disclose the information requested, or choose an “I prefer not to say” option (termed “active nondisclosure”). A further 6 items of a nonsensitive nature (e.g., season of birth) were included as filler items. A nondisclosure score was calculated by summing the number of items where an “I prefer not to say” option was chosen. This behavioral approach to the study of self-disclosure has been shown to be responsive to manipulations of privacy (Joinson, Paine, Buchanan, & Reips, 2008). “Passive” nondisclosure was not included in the count for two reasons: First, it is not possible to identify whether people are deliberately nondisclosing, so its inclusion may serve to conflate carelessness with privacy protection. Second, the rate of passive nondisclosure was too low for statistical analysis in its own right—on average, 4.5 participants left each self-disclosure item at the default (an average response rate of 0.59%).

Following the disclosure measures, participants completed measures of trust and perceived privacy designed to elicit their situational privacy and trust attitudes. Both measures were answered using a 5-point Likert scale (ranging from *strongly disagree* to *strongly agree*). The trust measure comprised eight items that incorporated the major dimensions of trust (Bhattacharjee, 2002; Jarvenpaa, Knoll, & Leiner, 1998; Metzger,

2004): *Benevolence* (e.g., “The intentions of this survey are good”; “The data I have provided will be kept secure and not exploited”), *Competence* (e.g., “This survey’s authors have the appropriate skills and competence to conduct online surveys”; “This survey is professional”), *Reliability* (e.g., “This survey’s authors are a dependable research group”), *Integrity* (“I do not doubt the honesty of this survey or its authors”; “The authors of the survey are trustworthy”), and *General trust* (e.g., “I felt comfortable giving my personal information”). Internal consistency reliability (alpha) for this measure was .91. Four additional filler items related to the design of the survey (e.g., “The design of the survey was clear”) and motivation (e.g., “I felt motivated to complete this survey”) were also included.

The perceived privacy measure had two questions relating to anonymity (“I felt anonymous completing this survey”) and confidentiality (“I am sure that my responses will remain confidential”), answered using the same scale. Alpha reliability for this measure was .73.

Procedure

An invitation to complete the study was sent to panel members by e-mail. For Time 1, members were informed that the survey consisted of a series of questions about any privacy concerns they may have when they use the Internet and their privacy related behavior. At Time 2, participants were told that some of the topics covered in the survey may be sensitive but that it was important for them to respond. The “prefer not to say” option was outlined, and they were told that the use of it would not imply any particular response.

At both time points, participants were informed that all information provided would remain confidential and that they could withdraw from the survey at any stage.

For all items participants were prompted to use the full scale and not only the labelled response options when responding. Participants’ responses were submitted and stored at the end of each page.

The Time 1 survey was left open for 2 weeks. Participants took, on average, 13 min to complete this part of the survey. Six weeks after data collection at Time 1 was complete, an invitation to complete the Time 2 survey was sent out to the same panel of participants. The delay between Time 1 and Time 2 was introduced to minimize the possible impact of the privacy measures on later disclosure behavior (Joinson et al., 2008). The Time 2 survey was left open for 2 weeks. Participants took, on average, 12.3 min to complete it.

2.2. Results

Descriptive Statistics

The mean scores, standard deviations, and Cronbach alphas for scales used are shown in Figure 1. The number of times an “I prefer not to say” option was chosen in response to a sensitive question was also summed to create a behavioral self-disclosure score, such that a higher score signified greater nondisclosure.

FIGURE 1. Descriptive Statistics and Reliability of Measures.

	Measure	Range	<i>M</i>	<i>SD</i>	Alpha
Privacy concern	Likert scale, high score = high privacy concern	19–80	58.90	12.10	.92
Privacy behavior: General caution	Likert scale, high score = more cautionary acts	10–30	21.70	5.21	.80
Privacy behavior: Technical protection	Likert scale, high score = more technical protection of privacy	6–30	22.61	5.59	.77
Nondisclosure	Sum of nominal responses (disclosed, active nondisclosure). High score = higher nondisclosure to items	0–10	0.45	1.05	.66
Trust	Likert scale, high score = high trust in survey group	10–40	32.10	5.69	.91
Perceived privacy	Likert scale, high score = high perceived privacy	2–10	7.66	1.90	.73

FIGURE 2. Correlations Between Measures.

	Privacy Concern	Privacy Behavior: General Caution	Privacy Behavior: Technical Protection	Perceived Privacy	Nondisclosure
Privacy behavior: General caution	.31**				
Privacy behavior: Technical protection	.15*	.29*			
Perceived privacy	-.04	.01	-.04		
Nondisclosure	.11**	.12*	.00	-.20	
Trust	-.05	.01	.04	.69**	-.21*

*Significant at the .05 level, two-tailed. **Significant at the .01 level, two-tailed.

Correlations Between Measures

Figure 2 shows the correlations between the various measures and the dependent variable (nondisclosure). To examine for potential confounding variables, a multivariate analysis of variance was calculated to examine any gender differences in responses to all the measures shown in Figure 1 (with the exception of nondisclosure). The multivariate tests found an overall effect of gender on these measures, $F(10, 653) = 8.84, p < .000, \eta^2 = .12$. Significant effects of gender were found for privacy protection: general caution ($p < .01, \eta^2 = .017, M = 22.25$ and 20.82 for female and male participants, respectively), privacy protection: technical protection ($p < .01, \eta^2 = .01, M = 23.50$ and 22.27 for female and male participants, respectively), and perceived privacy ($p < .05, \eta^2 = .007, M = 7.77$ and 7.44 for female and male participants, respectively).

Participants' age did not correlate with nondisclosure ($r_s = .024, p > .5$) but was related to privacy concern ($r_s = .10, p < .01$) and privacy behavior: general caution ($r_s = .14, p < .001$).

Privacy Attitudes, Reported Behaviors, and Self-Disclosure

A stepwise linear regression was calculated to examine the effect of the various privacy and dispositional variables on nondisclosure. Two further demographic variables—age and gender—were also added to the regression equation due to the significant associations just reported. The final step in the model is presented in Figure 3. In the first step, age and gender were entered (model $R^2 = .009$); in the second step, the dispositional measures (privacy concern, technical protection and general caution; model $R^2 = .03$); and in the final step, the two situational variables were entered (perceived privacy and trust, model $R^2 = .08$).

Model Testing

The regression analysis just presented provides evidence that privacy affects people's willingness to disclose information to a Web service in two forms: (a) through their dispositional privacy concerns and (b) through participants' perceived privacy during their interaction with the web survey (although this effect was only marginally significant). A further situational variable—trust—was the largest predictor. In the case of gender, female participants were less likely to disclose information than male participants.

In the following section, we focus on modelling the relationship between privacy concerns and situational variables in predicting nondisclosure. As noted in the introduction, others have argued that the impact of perceived privacy on behavior is influenced by trust (e.g., Metzger, 2004). However, as noted earlier, it is unclear if this relationship is one of mediation (wherein any relationship between perceived privacy and behavior is explained by common links to trust) or one of moderation (i.e., there is an interaction between perceived privacy and trust)

Furthermore, it is conceivable that the impact of dispositional privacy concern on disclosure behavior is also mediated or moderated by situational variables (i.e., per-

FIGURE 3. Predicting Nondisclosure: Privacy Concern, Trust and Demographics—Standardized Beta and t Values.

Step		Standardized β	t	Sig.
1	Gender	.09	2.50	.01**
	Age	.04	.94	.34
2	Gender	.07	1.92	.05*
	Age	.01	.27	.78
	Privacy concern	.08	2.16	.03*
	Privacy behavior: General caution	.09	2.20	.03*
	Privacy behavior: Technical protection	-.03	-.76	.44
3	Gender	.09	2.39	.02*
	Age	.02	.61	.54
	Privacy concern	.07	2.20	.03*
	Privacy behavior: General caution	.10	2.30	.02*
	Privacy behavior: Technical protection	-.04	-1.07	.29
	Overall trust	-.15	-2.76	.01**
	Perceived privacy	-.10	-1.89	.06

*Significant at the .05 level, two-tailed. **Significant at the .01 level, two-tailed.

ceived privacy or trust). If the relationship is mediated, then the relationship between privacy concern and nondisclosure is explicable through the links between privacy concern and the situational variables. However, a moderating relationship would suppose that there is an interaction between privacy concerns and situational variables to explain behavior.

Situational Factors and Nondisclosure


Baron and Kenny (1986) suggested that the test for a moderating relationship between variables should be through the creation of an interaction term and comparing the effect of this variable in a regression equation alongside the independent variables. This was done for the two situational variables in the present study by the creation of a multiplicative composite to examine the interaction between trust and perceived privacy. The first regression equation (perceived privacy and trust entered) explained 4.7% of the variance in nondisclosure. The addition of the interaction term in the second model increased prediction by 0.1% to 4.8%, but the interaction term did not significantly predict nondisclosure ($p = .18$). This suggests that there is not a moderating relationship between perceived privacy and trust on nondisclosure in the present study.

Mediation can be tested by examining the unique links between the two independent variables and the dependent variable (Baron & Kenny, 1986). To examine this possibility, regression equations are calculated between the two independent variables and between each independent variable and the dependent variable. To test for mediation, a series of regression models were calculated. In the first equation, the independent variable (IV; perceived privacy) significantly predicted the mediator (trust): Standardized $\beta = .71, p < .000$. In the second equation, the IV (perceived privacy) significantly predicted nondisclosure ($\beta = -.20, p < .000$). In the third equation, with both trust and perceived privacy entered, the mediator (trust) significantly predicted nondisclosure ($\beta = -.15, p < .01$ for trust; $\beta = -.09, p = .09$ for perceived privacy). The final test of mediation is that effect of the IV on the dependent variable (DV) must be lower in the third equation than in the second. This is indeed the case—the inclusion of trust as a mediator reduces the size of the standardized beta of perceived privacy by half. Because the effect of perceived privacy on nondisclosure is not reduced to zero, it can be noted that the effect of it on nondisclosure is partially mediated by trust. The results of these analyses are shown in Figure 4.

Situational and Dispositional Privacy and Trust and Nondisclosure

As noted in the introduction, it is also possible that the relationship between situational and dispositional privacy factors is also subject to either mediation or moderation. Again, to test for moderation, a multiplicative composite interaction term was created to test for the moderation of privacy concern by situational factors (perceived privacy and trust with the polarity reversed). In all cases, the variables were standardized to z scores. Again, regression equations were calculated to examine for any unique

FIGURE 4. Mediation of perceived privacy by trust on nondisclosure.



effect of this interaction term. In the first equation, the addition of an interaction term did not increase the amount of variance explained by privacy concern and perceived privacy alone as independent variables (5% vs. 4.9%).

In the second test of moderation (using trust) there was some slight evidence of an interaction between privacy concern and trust predicting nondisclosure. The first equation explained 5.8% of the variance in nondisclosure, whereas the introduction of the interaction term in the second equation increased this to 6.1%. This change in R^2 was approached significance ($p = .075$). In the first equation, both trust (Standardized $\beta = -.21, p < .001$) and privacy concern (Standardized $\beta = .12, p < .001$) were significant predictors. In the second equation, privacy concern remained significant (Standardized $\beta = .12, p < .001$), as did trust (Standardized $\beta = -.21, p < .001$), and the interaction term (Privacy Concern \times Trust) approached significance (Standardized $\beta = -.07, p = .075$). Thus, there is slight evidence that the interaction between trust and privacy concern is important in predicting nondisclosure, although the direct effect of privacy concern remained significant.

A potential mediator relationship between trust, privacy concern, and nondisclosure was also examined using the strategy just outlined. In the first equation, the IV (privacy concern) did not predict the proposed mediator (trust): Standardized $\beta = -.04, ns$. In the second equation, the IV (privacy concern) predicted nondisclosure (Standardized $\beta = -.11, p < .01$). In the third equation, with both trust and privacy concern entered, the proposed mediator (trust) significantly predicted nondisclosure (Standardized $\beta = -.21, p < .001$), as did privacy concern (Standardized $\beta = -.12, p < .01$). The final test of mediation is that effect of privacy concern on nondisclosure must be lower in the third equation than in the second. Clearly this is not the case, so there is no evidence of situational trust mediating the relationship between dispositional privacy concern and nondisclosure.

The same technique was also used to examine a possible mediator relationship between privacy concern, perceived privacy, and nondisclosure. Privacy concern did not predict perceived privacy (Standardized $\beta = -.04, ns$), and the inclusion of both IV and mediator did not diminish the effect of privacy concern on nondisclosure.

2.3. Discussion

As predicted, people's specific privacy concerns predicted their willingness to disclose personal information 6 weeks later. Two situational measures were also collected during the second part of the study: trust and perceived privacy. Both these measures also predicted people's willingness to disclose personal information to the Web site, although only marginally in the case of perceived privacy. Of importance, when entered into a regression equation, both dispositional privacy concerns (in the form of the Internet Privacy Concern scale) and situational factors (trust and perceived privacy) predicted disclosure to the Web site, suggesting independent effects for dispositional privacy concerns and situational perceived privacy and trust.

The results of the tests for moderation and mediation found evidence that, for situational aspects of privacy, the effect of perceived privacy on disclosure is mediated by trust. No such relationship was found between dispositional Internet privacy concern, trust, and disclosure, although there was slight evidence of a moderating relationship between trust and dispositional privacy concern on disclosure.

However, the evidence that the impact of perceived privacy on self-disclosure is mediated by trust should not be unequivocally accepted. In many cases, a moderator relationship is best tested using an experimental methodology (Baron & Kenny, 1986), especially when the nature of the relationship may be nonlinear. Moreover, the self-report measures used to assess situational trust and perceived privacy may be unduly influenced in Study 1 by the act of disclosure that preceded the completion of the measures.

The nature of the sample may also have influenced the results—the participants had an established relationship with the survey sponsor and may have been predisposed to trust the survey team. Gideon et al. (2006) reported that privacy attitudes had little impact on behavior when searching for a nonsensitive good, compared to a sensitive product. It is possible that sample characteristics (and, in particular, the relationship needed between investigators and respondents in longitudinal panel research) led to the low levels of variance in behavior explained by privacy concerns and levels of trust. It is also likely that distance education students who have signed up to an online panel are neither representative of the population from which they were drawn nor Internet users in general. This poses challenges for the generalizability of the results (something common in both online panels and longitudinal research) and suggests that the results should be examined using different populations.

The relatively low levels of nondisclosure in response to an "I prefer not to say" option has also been noted previously (Joinson et al., 2008). This pattern of varied privacy concerns, and almost universal disclosure, reflects similar results reported by Spiekermann et al. (2001), where privacy attitudes had little relationship to information divulged to a shopping 'bot. Although Joinson et al. (2008) provided evidence to support treating their behavioral nondisclosure measure as a scale (as used in Study 1, here), they also noted that there was evidence of a response set, such that participants tended to either disclose (or nondisclose) to all items, and exhibited little variance in their responses across the items. This may go some way, alongside the nature of the

sample, toward explaining the relatively small amounts of variance in self-disclosure explained by privacy concerns and trust. For this reason, fewer items were utilized in Study 2, and the behavioral self-disclosure measure was treated as a dichotomous variable.

Although other variables outside of privacy and trust no doubt influence willingness to disclose (e.g., conscientiousness), the relatively high homogeneity of the sample, the low level of nondisclosure across the sample, and their established relationship with the survey sponsor, may have contributed to the low levels of variance explained. The second study addresses these issues by recruiting participants via Internet Web sites, and by experimentally manipulating privacy and trust via different Web survey designs. Moreover, self-reports of perceived privacy and trust are also collected to enable a more in-depth examination of the interrelationship between the two variables at the situational level.

3. STUDY 2: EXPERIMENTAL MANIPULATION OF PRIVACY AND TRUST

In Study 2 privacy and trust were experimentally manipulated in a 2×2 between-subjects design using a Web-based survey. Participants were randomly allocated to one of four experimental groups (Privacy [High vs. Low] \times Trust [High vs. Low]), and their disclosure measured. Following the disclosure measures, participants completed measures of perceived privacy and trust in the survey process and authors.

3.1. Method

Participants

Participants were 181 Internet users recruited via advertisements on psychology and survey request Web sites. The majority ($n = 144$, 80%) were female (missing data for 5 people). Almost three fourths (73.9%) were based in the United States, with the remaining from 16 other countries (the majority from the United Kingdom and Canada). The age range spread from younger than 16 (1 person, removed from the analyses for ethical reasons) to older than 65 years, with the largest proportion aged 20 to 24 years (31.7 years).

Experimental Manipulations

In the privacy manipulations, the first page of the Web survey contained either a strong or weak privacy policy developed using the guidelines identified by Culnan (1999) and previously used in the field (e.g., Metzger, 2004; Miyazaki & Krishnamurthy, 2002; Nyshadham, 2000). Specifically, the strong privacy policy included information on the type of information collected, a notice that the information would not be reused or passed onto others, a list of security steps taken, and contact information. The weak privacy condition did not include full disclosure of information col-

lected and did not protect information from reuse or security lapses (see the appendices for the text of these statements). Pilot testing of the strong and weak privacy statement ($n = 57$) confirmed that the strong privacy policy was perceived as stronger at protecting privacy than the weak statement.

Trust was manipulated in a number of ways, based on the work of a number of researchers (e.g., Bhattacharjee, 2002; Fogg et al., 2001; Stanford, Tauber, Fogg, & Marable, 2002; Wang & Emurian, 2005). In the high trust conditions, the survey was hosted on an educational domain (*open.ac.uk), whereas in the low trust condition it was hosted on a domain designed to reduce trust (<http://www.surveylance.net>). The high trust condition included an institutional logo, no spelling mistakes, and no advertisements. The low trust condition incorporated advertisements (for gambling and money transfer services, links deliberately broken) and spelling and coding mistakes. Otherwise, the text within the Web pages was identical. Pilot testing ($n = 20$) confirmed that the trustworthy site was rated as significantly more trustworthy compared to the untrustworthy site ($p < .05$).

Measurement of Self-Disclosure

Disclosure was measured using the same technique outlined in Study 1. Participants completed four sensitive measures, each with an “I prefer not to say” option. The four items were those previously used in Study 1 with the highest levels of nondisclosure. These items were, “How many serious relationships have you had since age 18?” “How many sexual partners have you had?” “Are you a religious person?” and “What is your annual income?” The disclosure measures were followed by a series of demographic questions (age, gender, and country) alongside season of birth to maintain the face validity of the study.

Perceived Trust and Privacy

Perceived trust and privacy was measured using the same questions and response options as outlined in Study 1. A mean score for the trust items, and summed score for the perceived privacy items, was calculated for each participant.

Procedure

A link to the study was placed on a series of psychological and survey related Web sites (e.g., Web experimental lab). The study topic was advertised as “Life experiences and season of birth.” If participants clicked the link to the study, they were randomly allocated to one of the four conditions using JavaScript. A “no script” option directed them to a separate study. Only one participant was directed to this study using this link, suggesting that JavaScript was not an impediment to completion.

The experimental manipulation was embedded in the front page introducing the study. To proceed, participants clicked on a consent button and were then taken to a seriousness check; participants indicated on this page whether their answers should be included in the analyses. Following this, they then proceeded to the disclosure items (arranged on a single page). On submission, this page then led to the demographic

questions and finally the trust and perceived privacy items (introduced as a method to improve the quality of the surveys in the future).

3.2. Results: Self-Reported Trust, Perceived Privacy, and Experimental Condition

In previous research (e.g., Joinson et al., 2008; Joinson, Woodley, & Reips, 2007) the relatively low variance in nondisclosure has encouraged researchers to treat it as a dichotomous variable. Given the limited number of disclosure items in the present study, this approach was adopted here, with participant's responses dichotomized into those disclosing to all questions (76.1%) and those nondisclosing to at least one question (23.9%).

The proportion of nondisclosers and disclosers in each condition is shown in Figure 5. A chi-square test of association identified a significant association between condition and disclosure, specifically related to the combination of low privacy and low trust, $\chi^2(1, N = 95) = 7.28, p < .05$.

The pattern of results suggests an interaction between privacy and trust in determining people's willingness to disclose sensitive information. The combination of cues towards either high privacy or trust with cues toward low privacy or trust did not lead to reductions in disclosure behavior, suggesting that privacy and trust may operate in a compensatory manner. To examine this possibility, the impact of the experimental manipulations on participants' perceived privacy and trust were studied.

A mean score on the perceived trust items was calculated ($M = 3.88, SD = .76$) and a score on the two perceived privacy items calculated by summing the items ($M = 7.86, SD = 1.94$).

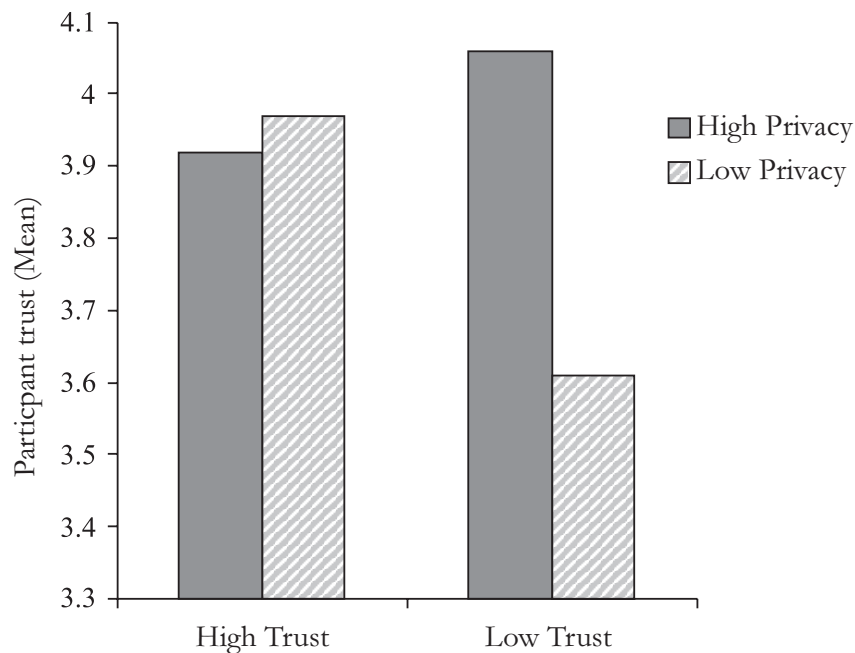
A two-way between-subjects analysis of variance (ANOVA) (Privacy \times Trust) found nonsignificant main effects of privacy condition, $F(1, 151) = 2.79, p = .09$, and trust condition, $F(1, 151) = 0.8, p = .37$, on participants' reported trust in the survey process and researchers. There was a significant interaction between privacy and trust condition on reported trust, $F(1, 151) = 2.44, p < .05$. This interaction is illustrated in Figure 6. The pattern of the interaction suggests that the inclusion of a strong privacy policy on the front page of the survey increased participant trust when combined with the low trustworthy design. Only when low trust was combined with a weak privacy policy did participant trust respond by falling substantially.

A second two-way between-subjects ANOVA (Privacy \times Trust) found a significant main effect of privacy condition, $F(1, 151) = 4.29, p < .05, \eta^2 = 0.028$, but not

FIGURE 5. Percentage of full disclosure by experimental condition.

	Privacy	
	High	Low
Trust		
High	78.3%	82.1%
Low	85.1%	60.4%

FIGURE 6. The interaction of privacy and trust condition on participant trust.



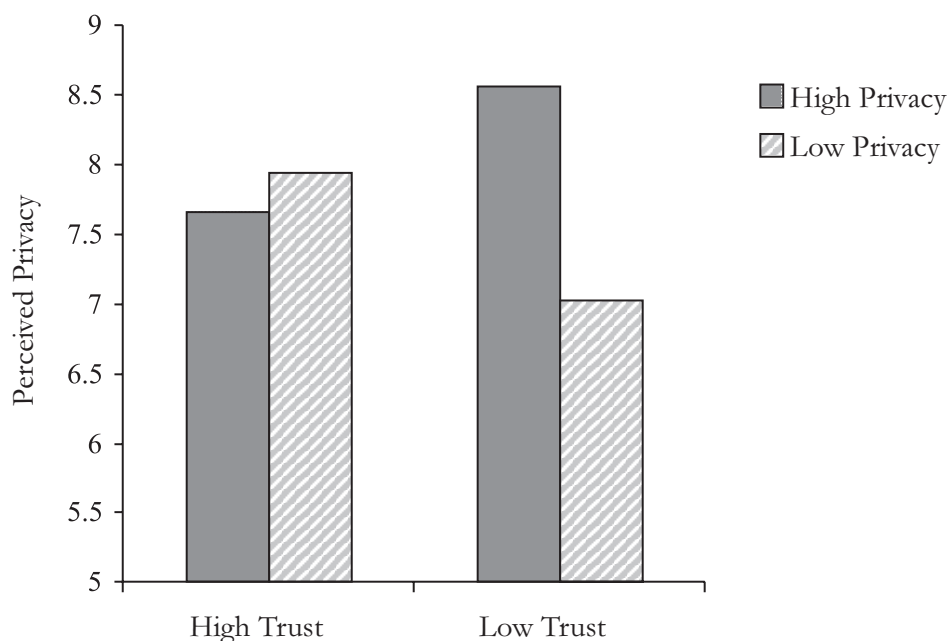
trust condition, $F(1, 151) = 0.01, p = .97$, on participants' reported perceived privacy. There was a significant interaction between privacy and trust condition on self-reported privacy, $F(1, 151) = 8.829, p < .01, \eta^2 = .055$. This interaction is illustrated in Figure 7.

The pattern of the interaction illustrates that the compensatory relationship between privacy and trust also exerts an influence on participants' perceived privacy. Participants tended to rate their perceived privacy relatively highly when a weak privacy statement was combined with cues designed to increase trust. However, the combination of a weak privacy policy and low trustworthiness led to lower levels of perceived privacy.

3.3. Discussion

The results of the present study demonstrate a strong moderator relationship between privacy and trust, as designed into a Web page, on both disclosure behavior and perceptions of trustworthiness and privacy. Self-disclosure was only substantially reduced when a weak privacy policy was combined with cues designed to reduce trust. In the conditions that combined high trust with low privacy, or low trust with high privacy, there was no evidence that self-disclosure was reduced. The second set of analyses explains this somewhat. A strong privacy statement, despite the presence of cues to lack of trustworthiness, increased participants' reported trust in the survey process and authors. Conversely, the presence of trust cues, when combined with a weak privacy statement, led to heightened perceived privacy. The compensatory interrelationship between trust and privacy is novel, and of particular value for understanding a large number of phenomena in the area of human–computer interaction (HCI) and privacy research. For instance, these results would suggest that studies that manipulate privacy but maintain

FIGURE 7. The interaction of privacy and trust condition on participant perceived privacy.



trust at a relatively high level may well report no significant association between privacy and behavior. Similarly, studies of trust that present a strong privacy policy may report no significant association, as high privacy compensates for low trust.

As outlined in the introduction (and replicated in Study 1), many researchers have proposed that the impact of privacy on behavior is mediated by trust. However, the results of Study 2 suggest that the relationship between privacy and trust may be significantly more nuanced than one of simple mediation. Specifically, the results suggest that the impact of privacy on behavior is *moderated* by trust, but that this moderation is not linear (and so would not be identified particularly well through the use of a multiplicative composite, as used in Study 1). The results also suggest that to fully understand people's reactions to potential privacy threats or actual violations it is imperative to also measure their trust in the privacy threat. For instance, people's privacy concerns over identity cards may only become salient or critical when combined with low trust in the government proposing such cards. Organizations that collect personal information are also, according to these results, unlikely to face substantial customer complaints until they lose trust. If, as suggested here, trust increases people's perceived privacy, this may go some way toward explaining why trusted organizations seem to be able to demand and collect vast quantities of information on consumers. The trust these organizations are recipients of may also increase the perceived privacy of their customers.

4. GENERAL DISCUSSION

The present studies are, to our knowledge, the first to include both situational and dispositional aspects of privacy and trust in the study of online disclosure, both

experimentally and survey based. It is important to note that in Study 1, we also separated the measures by 6 weeks, reducing the likelihood of any priming effect between the privacy measures and privacy related behavior, and measured actual behavior rather than reported actions or intentions.

The results of these two studies present strong evidence that privacy concern—both dispositional and related to the specific situation—influences people’s willingness to disclose personal information to a Web site. Although the effect was not large, it does suggest that when privacy concern is measured at an appropriate level of specificity, it does influence people’s behavior (cf. Ajzen & Fishbein, 1977). Second, the regression analyses in Study 1 suggest separate effects of a dispositional and situational privacy process that may have important implications for understanding the apparent disjuncture between people’s reported privacy concerns and their actual behavior. If there are no substantial links between people’s privacy concerns in general and their interpretation of the situation, then it would be expected that any link between general privacy concerns and behavior would be weak or nonexistent. The independent effects identified in Study 1 suggest that people’s interpretation of the trustworthiness of an organization, or their perceived privacy in a specific context, are not influenced by their general privacy disposition. Although many people may report high privacy concerns, when faced with a specific threat to their privacy in Study 1 they relied more heavily on situational cues to make a decision rather than their preexisting attitudes. In many ways, this reflects much of the ongoing work on the economics of privacy that stresses the costs and benefits to an Internet user of a specific situation requiring disclosure, rather than their disposition (e.g., Acquisti, 2005), as well as work that suggests that privacy is best understood as dynamic and active (e.g., Ahern et al., 2007; Palen & Dourish, 2003).

The results of Study 2 show that trust acts to moderate the impact of reduced privacy on both perceived privacy and behavior, whereas privacy also moderates the impact of reduced trust on behavior and reported trust. The symbiotic relationship between privacy and trust found in Study 2 goes a long way toward explaining why people may be willing to forgo privacy concerns when faced with a trusted requestor, and why privacy is important when faced with a request from an organization or individual one does not trust. It also supports previous research within the HCI literature that suggests that it is not only the nature of the information disclosed but also the recipient of the disclosure that is important (Olson et al., 2005).

5. LIMITATIONS AND FUTURE WORK

A number of limitations should be noted in the two studies reported in this article. First, in Study 1 the amount of variance in self-disclosure explained by the independent variables is relatively small. As noted earlier, this may in part be due to the nature of the measure and the relatively low variance in responses (self-disclosure was generally high across the board). Future work might wish to utilize other, more nuanced behavioral measures of disclosure, for instance, word count or content analy-

sis of disclosed information. Moreover, the nature of the sample in Study 1 may have limited nondisclosure—the respondents were distance education students who had previously agreed to complete a set number of surveys (i.e., they joined a panel). Quite aside from issues about the representativeness of the sample, this also raises issues in terms of their motivation, willingness to disclose information, and preexisting trust in the researchers. Study 2 addressed these concerns by the use of a different sample, but it is likely that nondisclosure was underrepresented in the first study, whereas levels of trust and privacy were higher than might be expected in a more naturalistic setting. Of course, it is difficult to conduct studies across time without building a relationship with participants, even if an identity management system were utilized. Moreover, in neither case were participants asked to disclose identifying information in the survey form, which may in itself have reduced the impact of privacy concern or trust on behavior. Future work might wish to include identifying information alongside the actual collection of sensitive responses to questions.

The relatively low amount of variance in self-disclosure explained by privacy and trust also highlights the importance of continuing research efforts to identify the factors that may influence people's decisions to disclose information to a Web site. For instance, there may be individual differences (e.g., propensity to openness, self-monitoring) that directly influence disclosure or impact upon levels of trust or privacy concern. The specific goals of the user may also impact on their decision-making process, as may the nature of the task conducted via the computer. Certainly, a strength of this article is that it has highlighted the complex and nuanced nature of the relationship between privacy, trust, and behavior while suggesting numerous future research programs based on the interaction between people's interpretation of a situation and their general preferences. For instance, the development of systems that seek to employ rules-based systems for the management of privacy may need to have designed in the opportunity for people to respond to specific circumstances on an ad hoc basis, as the results of the present research suggest that this situational path to determining privacy-related behavior is particularly important.

Finally, the results of the present study should be understood in terms of the potential limitations of the methodology and sample. Although considerable effort has been expended to create a controlled environment, the nature of this control may mean that the results do not generalize outside of the research context. Future research that studies these same relationships in naturalistic settings would be valuable.

6. CONCLUSIONS

Although many people express privacy-concerned attitudes, the link between these attitudes and their actual behavior has not been well established. In the present study, we demonstrated that, if an appropriate measure is used, it is possible to predict online behavior, albeit a relatively small amount of variance. We also found that two separate types of privacy-related attitudes—general and situational specific—pre-

dicted people's behavior. This has important implications for thinking about how we conceptualize users' privacy concerns and behavior. There was little evidence that people's dispositional privacy attitudes influenced their interpretation of the specific situation, suggesting that being concerned about privacy does not influence how a specific privacy-related situation is viewed. This could potentially offer an answer to why there seems to be a disjunction between privacy attitudes and behaviors, and certainly warrants further investigation.

In the second study, privacy and trust were found to operate in a symbiotic relationship, such that a lack of one was compensated for by a surfeit of the other. This insight has important implications for conceptualizing how organizations online manage to balance customer privacy requirements and trust, and illustrates the danger of neither protecting privacy nor engendering trust to online organizations.

NOTES

Background. An earlier version of the research reported in Study 1 was presented as a work-in-progress paper at CHI '06 (Paine et al., Privacy and Self-Disclosure Online, CHI '06 Extended Abstracts).

Acknowledgments. The authors are grateful to Jonathan Grudin and three anonymous reviewers for their insightful comments on this work.

Support. The work reported in this article was supported by funding from the UK Economic and Social Research Council E-Society Programme (RES-341-25-0011) awarded to Joinson, Buchanan, and Reips.

Authors' Present Addresses. Adam N. Joinson, School of Management, University of Bath, Bath, BA2 7AY, United Kingdom. E-mail: A.Joinson@bath.ac.uk. Ulf-Dietrich Reips, Departamento de Psicología, Universidad de Deusto Apartado 1, 48080 Bilbao, España. E-mail: reips@deusto.es. Tom Buchanan, Department of Psychology, University of Westminster, 309 Regent Street, London W1B 2UW, United Kingdom. E-mail: buchant@wmin.ac.uk. Carina Paine Schofield, Ashridge Public Leadership Centre, Ashridge Business School, Berkhamsted, Hertfordshire, HP4 1NS, United Kingdom. E-mail: Carina.Schofield@ashridge.org.uk.

HCI Editorial Record. First manuscript received May 2, 2007. Revisions received June 27, 2007 and March 18, 2008. Accepted by Jonathan Grudin. Final manuscript received May 1, 2009.—*Editor.*

REFERENCES

- Ackerman, M. S., & Cranor, L. F. (1999). Privacy critics: UI components to safeguard users' privacy. *Proceedings of CHI'99 Extended Abstracts*. New York: ACM.
- Ackerman, M. S., Cranor, L. F., & Reagle, J. (1999). Privacy in e-commerce: Examining user scenarios and privacy preferences. *Proceedings of ACM Conference on Electronic Commerce*. New York: ACM.
- Acquisti, A. (2004). Privacy in electronic commerce and the economics of immediate gratification. *Proceedings of the ACM Electronic Commerce Conference (EC '04)*. New York: ACM.

- Acquisti, A., Friedman, A., & Telang, R. (2006). Is there a cost to privacy breaches? An event study. *Proceedings of the International Conference of Information Systems (ICIS)*.
- Acquisti, A., & Grossklags, J. (2003). Losses, gains, and hyperbolic discounting: An experimental approach to information security attitudes and behavior. *2nd Annual Workshop on "Economics and Information Security"*.
- Acquisti, A., & Grossklags, J. (2005). Privacy and rationality in individual decision making. *IEEE Security and Privacy*, 26–33.
- Acquisti, A., & Grossklags, J. (2007, June). *When 25 cents is too much: An experiment on willingness-to-sell and willingness-to-protect personal information*. Paper presented at the 6th Annual Workshop on Economics and Information Security, WEIS 2007, Pittsburgh, Pennsylvania.
- Ahern, S., Eckles, D., Good, N. S., King, S., Naaman, M., & Nair, R. (2007). Photo sharing: Over-exposed?: Privacy patterns and considerations in online and mobile photo sharing. *Proceedings of CHI '07*. New York: ACM.
- Ajzen, I., & Fishbein, M. (1977). Attitude-behavior relations: A theoretical analysis and review of empirical research. *Psychological Bulletin*, 84, 888–918.
- Altman, I. (1975). *The environment and social behavior*. Monterey, CA: Brooks/Cole.
- Altman, I., & Taylor, D. (1973). *Social penetration: the development of interpersonal relationships*. New York: Holt, Rinehart, and Winston.
- Anderson, E., Charron, C., & Cohen, C.. (2005). The Consumer Privacy Bluff: Forrester Research. Available from: <http://www.forrester.com/Research/Document/Excerpt/0,7211,38299,00.html>
- Ang, L., & Lee, B.-C. (2000). Influencing perceptions of trustworthiness in Internet commerce: A rational choice framework. *Proceedings of Fifth COLLECTer Conference on Electronic Commerce*, 1–12. Brisbane, Australia.
- Baron, R. M., & Kenny, D. A. (1986). The moderator-mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations. *Journal of Personality and Social Psychology*, 51, 1173–1182.
- Bellotti, V., & Sellen, A. (1993). Design for privacy in ubiquitous computing environments. *Proceedings of the Third European Conference on Computer-Supported Cooperative Work 1993*, 77–92. Norwell, MA: Kluwer Academic Publishers.
- Ben-Ze'ev, A. (2003). Privacy, emotional closeness, and openness in Cyberspace. *Computers in Human Behavior*, 19, 451–467.
- Berendt, B., Gunther, O., & Spiekermann, S. (2005). Privacy in e-commerce: Stated preferences vs. actual behavior. *Communications of the ACM*, 48, 101–106.
- Bhattacharjee, A. (2002). Individual trust in online firms: Scale development and initial test. *Journal of Management Information Systems*, 19, 211–241.
- Buchanan, T., Paine, C., Joinson, A., & Reips, U.-D. (2007). Development of measures of online privacy concern and protection for use on the Internet. *Journal of the American Society for Information Science and Technology*, 58, 157–165.
- Chellappa, R. K., & Sin, R. G. (2005). Personalization versus privacy: An empirical examination of the online consumer's dilemma. *Information Technology and Management*, 6, 181–202.
- Cranor, L., Langheinrich, M., Marchiori, M., Presler-Marshall, M., & Reagle, J. (2002, April). *The Platform for Privacy Preferences 1.0 (P3P1.0) Specification*. Available from <http://www.w3.org/TR/P3P/>
- Culnan, M. J. (1999). *Georgetown Internet Privacy Policy Study: Privacy online in 1999: A report to the Federal Trade Commission*. Washington, DC: Georgetown University.

- Dutton, W. H., & Helsper, E. (2007). *Oxford Internet Survey 2007 Report: The Internet in Britain*. Oxford, UK: Oxford Internet Institute, University of Oxford. Available from: <http://www.oii.ox.ac.uk/microsites/oxis/publications.cfm>
- Fildes, J. (2006, May 24). Privacy worries over web's future. *BBC News*. Available from the BBC News Web site: <http://news.bbc.co.uk/1/hi/technology/5009774.stm>
- Fogg, B. J., Marshall, J., Laraki, O., Osipovich, A., Varma, C., Fang, N., et al. (2001). What makes web sites credible? A report on a large quantitative study. *Proceedings of CHI 2001*. New York: ACM.
- Galegher, J., Sproull, L., & Kiesler, S. (1998). Legitimacy, authority, and community in electronic support groups. *Written Communication, 15*, 493–530.
- Gefen, D. (2002). Reflections on the dimensions of trust and trustworthiness among online consumers. *ACM SIGMIS Database, 33*, 38–53.
- Gideon, J., Cranor, L., Egelman, S., Acquisti, A. (2006). Power strips, prophylactics, and privacy, oh my! *Proceedings of the Second Symposium on Usable Privacy and Security*.
- Harris, P. (2004). The European perspective—Is data protection value for money? *The 26th International Conference on Privacy and Personal Data Protection*.
- Harris and Associates Inc. (2004, June 10). *New national survey on consumer privacy attitudes to be released at Privacy & American Business Landmark Conference* [Privacy and American Business press release]. Available from http://www.marketwire.com/mw/release_html_b1?release_id=68484
- Jarvenpaa, S. L., Knoll, K., & Leidner, D. E. (1998). Is anybody out there? Antecedents of trust in global virtual teams. *Journal of Management Information Systems, 14*, 29–64.
- Jensen, C., & Potts, C. (2004). Privacy policies as decision-making tools: An evaluation of online privacy notices. *Proceedings of CHI '04*. New York: ACM.
- Joinson, A. N. (2001). Self-disclosure in computer-mediated communication: the role of self-awareness and visual anonymity. *European Journal of Social Psychology, 31*, 177–192.
- Joinson, A. N. (2008). “Looking at,” “Looking up” or “Keeping up with” people? Motives and uses of Facebook. *Proceedings of CHI 2008*. New York: ACM.
- Joinson, A. N., Paine, C. B., Buchanan, T., & Reips, U-D. (2008). Measuring self-disclosure online: Blurring and non-response to sensitive items in Web-based surveys. *Computers in Human Behavior, 24*, 2158–2171.
- 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.
- Jupiter Research. (2002a, May 20). *Security and privacy data*. Paper presented to the Federal Trade Commission Consumer Information Security Workshop. Retrieved June 20, 2005, from <http://www.ftc.gov/bcp/workshops/security/0205201leathern.pdf>
- Jupiter Research. (2002b, June 3). *Seventy percent of US consumers worry about online privacy, but few take protective action* (Survey summary). New York: Author. Retrieved from <http://web1.forrester.com/forr/reg/jupiterlogin.jsp>
- Malhotra, N. K., Kim, S. S., & Agarwal, J. (2004). Internet users' Information privacy concerns (IUIPC): The construct, the scale and a causal model. *Information Systems Research, 15*, 336–355.
- Margulis, S. T. (2003). On the status and contribution of Westin's and Altman's theories of privacy. *Journal of Social Issues, 59*, 411–429.
- Metzger, M. J. (2004). Privacy, trust and disclosure: Exploring barriers to electronic commerce. *Journal of Computer-Mediated Communication, 9*(4). Retrieved June 20, 2005, from <http://jcmc.indiana.edu/vol9/issue4/metzger.html>

- Metzger, M. J. (2006). Effects of site, vendor, and consumer characteristics on Web site trust and disclosure. *Communication Research*, 33, 155–179.
- Milne, G. R., & Culnan, M. J. (2004). Strategies for reducing online privacy risks: Why consumers read (or don't read) online privacy notices. *Journal of Interactive Marketing*, 18, 15–29.
- Miyazaki, A. D., & Krishnamurthy, S. (2002). Internet seals of approval: Effects on online privacy policies and consumer perceptions. *The Journal of Consumer Affairs*, 36, 28–49.
- Nyshadham, E. A. (2000). Privacy policies of air travel Web sites: A survey and analysis. *Journal of Air Transport Management*, 6, 143–152.
- Palen, L., & Dourish, P. (2003). Unpacking privacy for a networked world. *Proceedings of CHI '03*. New York: ACM.
- Pew Internet and American Life Project. (2001). *Trust and privacy online: Why Americans want to rewrite the rules*. Washington, DC: Author. Available from <http://www.pewinternet.org>
- Rubin, Z. (1975) Disclosing oneself to a stranger: Reciprocity and its limits. *Journal of Experimental Social Psychology*, 11, 233–260.
- Spiekermann, S., Grossklags, J., & Berendt, B. (2001). E-privacy in 2nd generation E-commerce: Privacy preferences versus actual behavior. *Proceedings of the Third ACM Conference on Electronic Commerce, Association for Computing Machinery (ACM EC'01)*, 38–47. New York, NY: ACM.
- Stanford, J., Tauber, E. R., Fogg, B. J., & Marable, L. (2002). *Experts vs. online consumers: A comparative credibility study of health and finance web sites*. Available from <http://www.consumerwebwatch.org/news/report3credibilityresearch/slicedbreadabstract.htm>
- Tidwell, L. C., & Walther, J. B. (2002). Computer-mediated communication effects on disclosure, impressions, and interpersonal evaluations: Getting to know one another a bit at a time. *Human Communication Research*, 28, 317–348.
- Tsai, J., Egelman, S., Cranor, L., & Acquisti, A. (2007). *The effect of online privacy information on purchasing behavior: An experimental study*. Paper presented at the 6th Annual Workshop on Economics and Information Security (WEIS 2007), Pittsburgh, Pennsylvania, 7-8 June 2008.
- UK Information Commissioner. (2006). *Report on the surveillance society*. Retrieved November 27, 2006, from <http://tinyurl.com/ya76db>
- UPI-Zogby. (2007, April 3). *Zogby poll: Most Americans worry about identity theft*. Available from <http://www.zogby.com/NEWS/ReadNews.dbm?ID=1275>
- Wang, Y. D., & Emurian, H. H. (2005). An overview of online trust: Concepts, elements, and implications. *Computers in Human Behavior*, 21, 105–125.

APPENDIX A1. HIGH PRIVACY FRONT PAGE

The information that you choose to provide during this survey will be collected and stored by the researchers, Dr. Adam Joinson [mailto: tag]. and Dr. Carina Paine [mailto: tag], within a secure database at the Open University, a large higher education institution in the United Kingdom.

You can choose not to provide information to any item, and you are free to withdraw from the survey at any stage by simply closing the window. If you do withdraw

from the survey, any information provided will be discarded from the data file. We do not ask for, or collect, any information that would enable you to be personally identifiable, and any analysis is based on the combined responses of all respondents, not individuals.

Any information you provide will remain strictly confidential - we will not share the responses you give in the study with anyone else outside of the research team. All information provided will be stored in a secure location, and will not be taken outside of the Open University on portable media.

Please feel free to contact us if you have questions or concerns. Our email address is Elsa-Presto@open.ac.uk. You can contact us by post using the following address: Institute of Educational Technology, The Open University, Milton Keynes, MK7 6AA.

If you are happy to continue please click below to go to the first page of the survey. If you have any questions, please e-mail us using this address [[mailto: tag](mailto:tag)].

APPENDIX A2. LOW PRIVACY FRONT PAGE STATEMENT

The information that you choose to provide during this survey will be collected and stored by the researchers at the Open University, a large higher education institution in the United Kingdom. Backup copies of the datafile will be kept off-site on portable media (memory sticks and CD-ROM). We may also pass your information on to our collaborators for further analysis.

You can choose not to provide information to any item, and you are free to withdraw from the survey at any stage by simply closing the window.

While we will do all we can to protect your privacy, you should be aware that your responses could be intercepted by third parties such as hackers or law enforcement agencies.

Your browser may also store data locally in your computer's hard drive, with the result that other people using your computer could potentially find it. The data you submit may also be held in a temporary store or cache maintained by your university, Internet service provider or employer.

Web-servers automatically collect i.p. numbers for each person who visits a page. These numbers are the 'Internet address' of the computer you are using. We use the i.p numbers associated with each response to check for people completing the survey more than once.

If you are happy to continue please click below to go to the first page of the survey. If you have any questions, please e-mail us using [[mailto: tag](mailto:tag)].