Measuring self-disclosure online: Blurring and non-response to sensitive items in web-based surveys

Adam N. Joinson [a,*], Carina Paine [b], Tom Buchanan [c], Ulf-Dietrich Reips [d]

[a] *School of Management, University of Bath, Bath BA2 7AY, UK*
[b] *Institute of Educational Technology, The Open University, MK7 6AA, UK*
[c] *Department of Psychology, University of Westminster, UK*
[d] *Department of Psychology, University of Zurich, Switzerland*

Abstract

People are increasingly required to disclose personal information to computer- and Internet-based systems in order to register, identify themselves or simply for the system to work as designed. In the present paper, we outline two different methods to easily measure people's behavioral self-disclosure to web-based forms. The first, the use of an 'I prefer not to say' option to sensitive questions is shown to be responsive to the manipulation of level of privacy concern by increasing the salience of privacy issues, and to experimental manipulations of privacy. The second, blurring or increased ambiguity was used primarily by males in response to an income question in a high privacy condition. Implications for the study of self-disclosure in human–computer interaction and web-based research are discussed.

1. Introduction

Self-disclosure is the telling of the previously unknown so that it becomes shared knowledge, the "process of making the self known to others" (Jourard & Lasakow, 1958, p. 91).

---

[*] Corresponding author. Tel.: +44 1225 386233; fax: +44 1225 386473.
E-mail address: A.Joinson@bath.ac.uk (A.N. Joinson).

This shared knowledge might exist between pairs of people, within groups, or between an individual and an organization. It has a variety of purposes, in part dependent on the context in which disclosure occurs. For instance, within dyads, particularly romantic relationships, it serves to increase mutual understanding (Laurenceau, Barrett, & Pietromonaco, 1998), and builds trust by making the discloser increasingly vulnerable (emotionally or otherwise) to the other person (Rubin, 1975). Since self-disclosure is often reciprocated within a dyad, it often serves to strengthen the ties that bind people in romantic or friendship-based relationships (Jourard, 1971). Disclosure within groups can serve to enhance the bonds of trust between group members, and it can also serve to legitimize group membership and strengthen group identity (Galegher, Sproull, & Kiesler, 1998). Finally, disclosure between an individual and an organization can serve authentication purposes – for instance, to establish identity, allow authentication of a claim to identity and to enable an organization to recognize you in the future in order to personalize its offerings to you. Organizations might also ask for personal information for marketing purposes – for instance, when registering to access a website or joining an online community. Of course, organizations, in the form of researchers, might also ask for personal information in the name of academic or commercial research.

Self-disclosure is of particular interest to researchers of computer-mediated communication and human–computer interaction for a number of reasons. First, since it acts as a signifier of trust and acceptance of privacy assurances, it is of interest in terms of e-commerce and relationship building online (e.g. Bargh, McKenna, & Fitzsimons, 2002; Joinson & Paine, 2007). Second, self-disclosure is critical for many Internet-based services (e.g. personalization, Web 2.0 software) to function effectively (Chellappa & Sin, 2005) or for Internet applications (such as e-recruitment; Nickel & Schaumburg, 2004) to be used to their fullest potential. In the context of person-to-person mediated interaction, self-disclosure has been found to be higher online compared to face-to-face, as well as being responsive to changes in the design of interaction systems (Joinson, 2001a; Joinson & Paine, 2007; Tidwell & Walther, 2002). For researchers, there is also considerable evidence that computer-mediated and web-based survey hosting can encourage increased disclosure and candid responses (Joinson, 1999; Tourangeau, 2004). Explanations for heightened disclosure online usually refer to the reduced vulnerability that online anonymity affords the discloser (the 'Strangers on the Internet' phenomenon). Psychologically, anonymity has also associated with increased self-disclosure via changes in self-awareness processes (Joinson, 2001a), uncertainty seeking behavior (Tidwell & Walther, 2002) and an online disinhibition effect (Suler, 2004).

However, the use of self-disclosure in web-based studies has been hampered by the measures available, and specifically the time taken to analyze open-ended responses or interactions. For instance, within computer-mediated communication research, it is common to utilize some form of content analysis, either to score the breadth of disclosure using word-count, or the depth of such disclosure using rating systems (e.g. Joinson, 2001b). This approach has also been applied in studies of human–computer interaction, for instance, Spiekermann, Grossklags, and Berendt (2001) used an automated, bot to engage in question and answer sessions with users and scored both breadth and depth of disclosure. However, this technique is time consuming, and is not applicable to many experimental studies where such discourse based measures do not fit easily with experimental manipulations. A further complication is that not all self-disclosure is equal – disclosing your season of birth is not the same as disclosing your age, which is not the same as disclosing your sexual fantasies. According to Altman and Taylor (1973), self-disclosure varies according
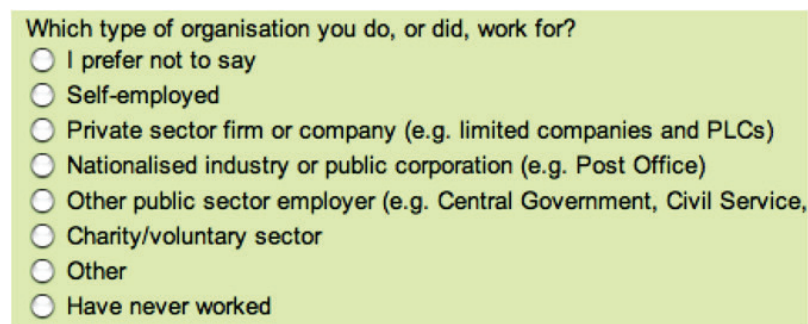
to the *breadth* (i.e. amount of information disclosed), *depth* (i.e. the intimacy of that information) and *duration* (i.e. the amount of time spent disclosing). In terms of the depth of disclosure, they argue that disclosed information can reveal either peripheral, intermediate, and core layers of the self. The peripheral layer is concerned with biographic data (e.g. age), clothing and so on. The intermediate layer deals with attitudes and opinions and the core layer with personal beliefs, needs, fears, and values. However, as Tidwell and Walther (2002) note, such classification schemes pose a number of practical issues, including overlapping boundaries (e.g. sexual preference is both biographic and core), while topics could potentially overlap across a number of different layers. Antaki, Barnes, and Leudar (2005) also argue that the act of disclosure needs to take into account the interactional context rather than simply be scored on a checklist. For instance, the phrase "I'm the worlds worst cook" could be disclosure, a plea for help or self-deprecation. Without the context, argue Antaki et al., it is not possible to be certain.

Alternatively, self-report measures of disclosure have been used with some success, for instance to compare levels of disclosure in face-to-face (FtF) and online relationships, or to link marital satisfaction with disclosure within the relationship. For instance, Parks and Floyd (1996) asked their participants to report the level of self-disclosure in their Internet relationships using self-report (e.g. high scores on "I usually tell this person exactly how I feel" and low scores on "I would never tell this person anything intimate or personal about myself"). However, the same problems – a lack of context – arise for such self-report measures too.

In the present paper we present two alternative methods for studying self-disclosure that can be used in web-based studies. The first measure relies on the provision of a response option to not disclose to sensitive personal questions. There are two main ways in which non-response can be operationalized. The first is either submitting a default selection, or where there is no default option, submitting no response. A second is to add an option that allows participants to select 'I prefer not to answer' (Buchanan, Joinson, & Ali, 2002; Knapp & Kirk, 2003). If this option is visible on a web page, but is not the default, then it is possible to distinguish non-response (use of defaults or no answer) from purposeful secrecy or non-disclosure (see Fig. 1).

Importantly, this technique can also be easily integrated into a number of different scenarios; for instance, we have used it as part of a registration form, and embedded it within the demographic section of a web-based questionnaire. The problems with context noted above are eliminated in this technique, because disclosure is actually being measured within the context of interest.

While it has been argued that the provision of 'no opinion' choices may increase satisfying in attitude surveys (Holbrook, Green, & Krosnick, 2003), there is little reason to

Fig. 1. Item with 'I prefer not to say' option.

What is your personal annual income, before taxes
Between [ ] and [ ] pounds
○ I prefer not to say

Fig. 2. Example of range item to measure blurring.

assume that a similar process would operate in the use of 'I prefer not to answer' responses to sensitive personal questions. Indeed, Joinson, Woodley, and Reips (2007) report that the provision of 'I prefer not to answer' options in a salary question may improve data quality by reducing the number of non-responses or default selections.

The second technique we study in the present paper is 'blurring'. Ambiguity is an established method for the disclosure of non-diagnostic information in location-based systems (Duckham & Kulik, 2005; Gruteser & Grunwald, 2003). For instance, in response to the question, "where are you?" the answer 'in London' provides information of lower diagnostic value compared to, say, "Starbucks on the Tottenham Court Road". In the present paper, we examine the potential use of blurring as another technique to study self-disclosure. The advantage of blurring in many human–computer interactions is that it provides an alternative option compared to simply disclosing or not, and is therefore potentially more powerful for the study of nuances in disclosure rather than simply treating disclosure as a nominal variable. It could also be argued that while the use of an 'I prefer not to say' option studies the breadth of disclosure, the provision of blurring opportunities provides a mechanism for the study of the depth of disclosure (Altman & Taylor, 1973, discussed above). In the present study, blurring is operationalized through the use of open-ended range items (see Fig. 2). A large range (e.g. $10,000–$100,000) would indicate blurring, compared to a smaller range (e.g. $10,000–$12,000).

In the present paper, we outline two studies that examine these techniques for use in the study of self-disclosure in web-based studies. In the first study, a series of items with 'I prefer not to say' are tested. The sensitivity of these items to changes in privacy salience is tested through the use of an experimental manipulation designed to prime privacy concerns. Half the participants complete a measure of Internet privacy concern before the self-disclosure items (the privacy prime group), the remainder complete the self-disclosure items first, followed by the privacy concern measure (the non-primed group). It is predicted that, if people's responses to the items are sensitive to changes in privacy salience, the privacy prime group would be more likely to utilize 'I prefer not to say' response options compared to the non-prime group. No specific prediction is made for non-sensitive items since response patterns to both sensitive and non-sensitive items may not be independent.

In the second study, use of 'I prefer not to say' and blurring are considered in more detail in light of both gender and privacy, two variables previously identified as having an effect on self-disclosure (Dindia & Allen, 1992; Joinson & Paine, 2007).

2. Study 1

2.1. Participants

To examine the likelihood that 'I prefer not to say' responses would be used, 515 members of a volunteer research panel of Open University (OU) students completed a series of

measures, including the self-disclosure items. In total 685 members of the panel were invited by e-mail to complete the web-based questionnaire (response rate: 75%). Of the 515 respondents, 43% (220) were male and 57% (286) were female (demographic data was unavailable for nine participants). The mean age of the sample was 43.9 years (range: 22–77 years, SD = 10.4).

2.2. Measures and procedure

Participants completed an 18-item measure of self-disclosure titled 'About you and your life experiences'. Six of the items were non-sensitive (e.g. "What season were you born in", "Are you left or right handed", and "What is your height"?) and twelve sensitive (e.g. "How many sexual partners have you had", "Since age 18 how many different serious relationships have you had" and "Are you a religious person"?). In total, 30 self-disclosure items were developed.[1] The items had been subjected to pre-testing to identify sensitive and non-sensitive items based on people's reported willingness to disclose the information requested (Buchanan et al., 2002). To reduce participant effort, the sensitive items were grouped for face validity and split into two 12-item versions. The six non-sensitive items were presented in both versions. Participants responded to the items using both drop down menus and open text boxes (for instance, when height and weight were asked). For drop down responses, the default option was "Please select" with "I prefer not to say" the second choice. For text boxes, a radio button labeled "I prefer not to say" was provided next to the text box. Participants were instructed:

> We appreciate that some of the topics covered may be sensitive and have therefore provided a "prefer not to say" response option for all of the questions. Please feel free to use this option for any item and be assured that this does not imply any particular response.

Participants also completed a measure of privacy concern and behavior (Buchanan, Paine, Joinson, & Reips, 2007). This measure comprises 16 Internet privacy attitude items (e.g. "Are you concerned about people you do not know obtaining personal information about you from your online activities?" and "Are you concerned about online organizations not being who they claim they are?") and 12 privacy behavior items (e.g. "Do you clear your browser history regularly?" "Do you read a website's privacy policy before you register your information?"). The items were answered using a 5-point Likert scale. The privacy measure was used as a prime for non-disclosure, such that half the participants completed the self-disclosure measure first, then the privacy measure. The pattern was reversed for the other half. It is considered that completing the privacy questionnaire is likely to increase the salience of privacy issues to participants. We expect that among participants for whom privacy concerns have been so primed; non-disclosure is likely to be higher.

2.3. Results and discussion

To enable more detailed analyses, and to match the number of non-sensitive items, the six sensitive items with the highest percentage of people selecting 'I prefer not to say' were

---

[1] A copy of the full set of self-disclosure items and response options is available by e-mailing the first author.

Table 1
Percentage use of 'I prefer not to say" for six most common sensitive items, and the non-sensitive items

	% Selecting "I prefer not to say"	% Non-response
Six most frequently non-disclosed items		
How many different sexual partners have you had?	27.30	0.39
Since age 18 how many different serious relationships have you had?	14.04	0.78
Have your partners been [all same gender, all opposite gender, a mixture]	13.26	1.15
What is your weight?	12.80	0.00
How long is/was your longest relationship to date?	10.25	0.39
Are you a religious person?	9.92	0.78
Non-sensitive items (selected from pre-testing)		
Are you a religious person?	9.92	0.78
Do you carry a donor card?	3.10	0.58
Do you give to charity?	3.31	1.94
Are you left or right handed?	1.56	0.69
Which season were you born in?	1.56	0.39
What is your height?	1.12	0.00

identified (see Table 1). Non-response submissions (i.e. no option selected) are also presented in the table. One non-sensitive item (religiosity) had a relatively large percentage of people choosing the 'I prefer not to say' option (around 10%). For this reason, this item is excluded from the following analyses, and is treated as a sensitive item in Study 2.

Participants' responses to the five remaining sensitive items were summed to create a non-self-disclosure score. Use of 'I prefer not to say' was scored as '1', and a response to the item scored as '0'. The same procedure was used to create a non-self-disclosure score for the sensitive items. The resulting non-disclosure score was significantly higher for sensitive items (Mean = 0.8, SD = 1.34) compared to the non-sensitive items (Mean = 0.12, SD = 0.48, paired $t(256) = 9.02$, $p < 0.001$).

Two one-way ANOVAs were computed to compare non-disclosure rates in light of the privacy priming manipulation. In the first ANOVA, non-disclosure to the sensitive items was significantly higher in the privacy-primed condition (Mean = 0.96, SD = 1.47) compared to the non-primed condition (Mean = 0.62, SD = 1.21, $F(1, 255) = 4.16$, $p < 0.05$). In the second ANOVA, use of the 'I prefer not to say' option to non-sensitive questions was also significantly higher in the privacy prime condition (Mean = 0.2, SD = 0.62) compared to the non-prime condition (Mean = 0.05, SD = 0.28, $F(1, 255) = 6.70$, $p < 0.01$). This pattern of results shows that the use of an 'I prefer not to say' option to measure self-disclosure is responsive to changes in the salience of privacy concerns, and that this effect occurs for both responsive and non-sensitive questions.

The results of the first study suggest that the use of a series of sensitive questions with 'I prefer not to say' options may be feasible as a measure of self-disclosure/secrecy. This is justified for a number of reasons. First, the use of 'I prefer not to say' was significantly higher for sensitive compared to non-sensitive questions. Second, the use of 'I prefer not to say' increased following the priming of privacy concerns, compared to those who were not primed with these concerns. However, the effect of priming applied regardless of the sensitivity of the questions. This could be for a number of reasons. First, despite pre-testing, the non-sensitive questions used in the present study may not have been as innocuous as expected. This particularly seems to be the case for religiosity, and possibly

for possession of a donor card and donation to charity. Second, it is possible that self-disclosure to a single item is not independent of the context, and a response set may be established among participants such that non-disclosure to sensitive items encourages later non-disclosure to non-sensitive items. Future research should address the possible inter-dependence of both sensitive and non-sensitive items. To partially address this possibility, in the second study the use of 'I prefer not to say' is treated as a dichotomous variable, rather than as continuous. Finally, for some of the items in Study 1, the use of an 'I prefer not to say' option could potentially be seen as diagnostic in its own right. For instance, refusal to answer a question about, for instance, drug use, may lead the questioner to interpret refusal as an affirmative answer. This possibility is examined in Study 2 through the examination of both non-disclosure and the provision of alternative methods to disclose while protecting privacy. As noted in the introduction, self-disclosure can be understood in terms of both the breadth and depth. The use of an 'I prefer not to say' option addresses the breadth of the disclosure, but not the depth. In the second study, a method to study the depth of disclosure – called 'blurring' is studied.

3. Study 2

The aim of Study 2 was to examine the utility of the 'I prefer not to say' method in a more naturalistic setting. In this study, sensitive items (income, religiosity and ethnicity) are embedded in a registration form. Other methods to measure disclosure and secrecy (i.e. range items to enable blurring and deception checks) are also included in the study. Rather than priming privacy (as used in Study 1), privacy is experimentally manipulated in Study 2.

3.1. Overview

The context for this study was a request sent to part-time adult students to register for an online opinion survey panel. The sample was adult distance education students who commonly combine work and/or childcare responsibilities with part-time study. A total of 10,000 students were sent an email from the University Principal asking if they would consider signing up to the panel for a period of 12 months. The panel sign-up process used in the present study occurs every 12 months. The items used in the present study are those normally collected during the process of registration, with the addition of an 'I prefer not to say' option and blurring opportunities the only deviation from the normal registration form. Participants are made aware that the data collected is subject to the Data Protection Act (1998, UK), and that any responses are treated confidentially, will not be released, and will not form part of their student record. The study was given ethics approval by the Open University Human Participants and Materials Ethics Committee (Ref: HPMEC/07/#338/).

If participants clicked on the link they were directed to a registration website comprising three pages: a front page outlining the panel, a sign-up page that asked for their forename, surname, email address and personal identifier number (assigned on registration with the University), and a demographic questions page asking a series of standard demographic queries. The self-disclosure measures were embedded within this demographic page.

The impression given to participants of their privacy was experimentally manipulated using a number of well-established methods (Culnan, 1999; Joinson et al., 2007). In the low privacy condition, participants were (a) sent a personalized e-mail, (b) directed to a

front page without a privacy statement (it instead outlined the project with matched word count), and (c) found that their personal information had already been loaded into the sign-up page. In the high privacy condition, participants (a) were sent a non-personalized e-mail, (b) were directed to a front page with a strong privacy statement, and (c) did not have their personally identifiable information pre-loaded. The strong privacy policy was developed using the guidelines identified by Culnan (1999) and previously used successfully (e.g. Metzger, 2004; Miyazaki & Krishnamurthy, 2002; Nyshadham, 2000). Specifically, the strong privacy policy included information on the type of information collected; reassurance that it would not be re-used or passed onto others, the security steps taken and contact information for those collecting the information. We also counterbalanced the order of the sign-up page requesting personally identifiable information, and the demographic questions page.

3.2. Participants

Participants were 1189 current Open University students who responded to a request to join an online research participant's panel. There were 748 females (62.9%) and 441 males (37.1%) in the sample. The mean age of the participants was 42.75 years (SD = 11.81).

3.3. Measures

A series of self-disclosure measures were embedded within the demographic questions:

3.3.1. Deception
To examine the possibility of deception as a strategy of non-disclosure, participants completed a question where verifiable information was already available – date of birth (dd/mm/yy).

3.3.2. I prefer not to say
While an 'I prefer not to say' option was provided for all questions (except the open text boxes), we focused in particular on three items that we had established in pre-testing as sensitive that also retained face validity in the specific context (a registration form). These were annual salary, religiosity and ethnic group.

3.3.3. Blurring
A number of items were provided with open-ended ranges. For instance, salary could either be non-disclosed using 'I prefer not to say' or responded to by typing top and bottom ranges into two separate text boxes (see Fig. 2). Thus, participants had the opportunity to disclose with high informational value (salary $25,000–$25,100) or low informational value (salary $10,000–$70,000). Range items using this technique were income, hours spent watching TV, hours spent using the Internet, novels read per year and hours spent studying (per week).

3.4. Results and discussion

More people completed the sign-up process in the low privacy condition (n = 722) compared to the high privacy condition (n = 467). The effect of personalized email salutation

in increasing response rates has been previously documented (Joinson & Reips, 2007). We also suspect that the reduced effort required (e.g. no need to seek out personal identifier code) also contributed to this increased response rate.

3.4.1. Deception as a privacy protecting strategy

The vast majority of respondents ($n = 1117$, 95%) disclosed their date of birth. Of these, the majority (90%) disclosed a date of birth that matched the one recorded on their student record. Unfortunately 9.1% of students entered "19" as the year of their birth. This was perhaps due to the fact that the text field for the year of birth was only two characters wide. The day and month of birth for these students was entered correctly.

Six students (0.5%) made what appeared to be "genuine" mistakes (e.g. incorrectly entering 18/4/43 instead of 18/1/43) or there may have been mistakes in the original student records (e.g. 19/7/80 instead of 19/7/81). Finally, it appeared that nine respondents "deceived" by providing a false date of birth. Two of these students deceived on two parts of their response (day, month or year) and the remaining seven students deceived on all parts of their response. This is a low deception rate (less than 1%).

3.4.2. Use of 'I prefer not to say'

Income was the item most commonly non-disclosed to (37.5% of responses were 'I prefer not to say', see Table 2). This was followed by religiosity (10.7% non-disclosure) and ethnicity (3.9% non-disclosure).

Chi square tests showed a significant association between gender and non-disclosure to both the income and ethnicity questions (see Table 2), specifically that women were less likely to disclose their income than men, and men were less likely to disclose their ethnicity than women.

The three sensitive items (salary, religiosity and ethnicity) were combined to create a single non-disclosure score, such that non-disclosing to all items scored 3, and disclosing to all items scored 0. The mean score on this measure was 0.50 (SD = 0.66). Only 11 people chose 'I prefer not to say' to all three questions. This limited variance in disclosure, and possible floor effect, suggests that the variable may be better treated as dichotomous rather than as continuous. Thus, the self-disclosure variable was dichotomized into those disclosing to the majority of questions (1102), versus those non-disclosing to the majority of questions (87). A chi-square test of association confirmed that significantly more people used 'I prefer not to say' to the majority of their answers in the low privacy condition (8.7%) compared to the high privacy condition (5.1%; $\chi^2(1, 1189) = 5.38$, $p < 0.02$; odds ratio = 1.76).

Men and women showed a different pattern of non-disclosure to the privacy manipulations (see Table 3). While men tended to disclose more than women (see Table 2 below), for two of the items they showed responsiveness to the privacy manipulation such that disclosure was reduced in the low privacy condition. For females, the level of disclosure

Table 2
Non-disclosure (percentages) by sensitive item and sex

	Male (%)	Female (%)	χ^2 (df, n)
Income	32.5	40.6	7.16 (1, 1140), $p < 0.01$
Religion	9.1	11.6	1.54 (1, 1159), ns
Ethnicity	5.6	3.0	4.09 (1, 1166), $p < 0.05$

Table 3
Percentage use of 'I prefer not to say' by privacy condition and gender

	Females			Males		
	High privacy (%)	Low privacy (%)	χ^2 (df, n)	High privacy (%)	Low privacy (%)	χ^2 (df, n)
Income	42.6	39.4	0.7 (1, 712), ns	28.3	35.9	2.70 (1, 428), $p = 0.06$
Religion	11.1	11.9	0.12 (1,731), ns	6.3	11.3	3.23 (1, 428), $p = 0.05$
Ethnicity	1.1	4.0	3.99 (1, 735), $p < 0.05$	4.1	6.7	0.91 (1, 431), ns

remained stable across the two privacy conditions for income and religiosity, while their use of the 'I prefer not to say' option was significantly higher (albeit low in both conditions) in response to the ethnicity question in the low privacy condition.

Standardized (z) scores were calculated for the range values provided for income in order to compare the degree of blurring by males and females. The mean difference between the two range values for females was −0.02 (SD = 0.19), indicating a score below the mean. For males, the mean difference was above the mean (mean difference = 0.03; SD = 0.29). Therefore, although a significantly higher percentage of males disclosed their income (see Table 2), they "blurred" their actual income details by disclosing a larger range. Table 4 shows the mean differences for males and females in each privacy condition (using standardized scores).

A two-way ANOVA (gender × privacy condition) with the standardized scores for the degree of blurring as the dependent variable showed a significant main effect of gender [$F(1, 341) = 3.884$, $p < 0.05$, partial $\eta^2 = 0.011$] and a significant privacy × gender interaction [$F(1, 341) = 5.140$, $p < 0.05$, partial $\eta^2 = 0.015$], which is illustrated in Fig. 3. This interaction shows that females were always below mean in terms of blurring income, while

Table 4
Blurring by gender and privacy condition

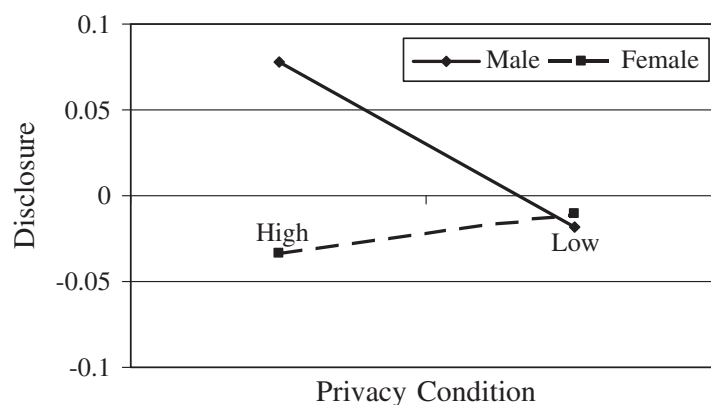| Mean difference (SD)			
Male		Female	
High privacy	Low privacy	High privacy	Low privacy
0.078 (0.308)	−0.019 (0.256)	−0.034 (0.144)	−0.011 (0.207)

Fig. 3. Gender, privacy and blurring (standardized scores, mean = 0).

males were above the mean in the high privacy condition, and below the mean in the low privacy condition.

This pattern of results suggests that participants' privacy condition has a differential impact on males and females' privacy-enhancing responses, such that females non-disclosed their salary equally in both the high and low privacy condition, while males non-disclosed in a low privacy condition, and disclosed but blurred in a high privacy condition.

4. Conclusions

The measurement of self-disclosure in web-based studies is often time consuming or problematic. In this paper, we have outlined a number of ways in which disclosure, and its converse, secrecy, can be easily measured using simple web-based forms. Importantly, these measures have the face validity to be of use in a number of different scenarios, including for instance, registration forms and basic demographic questions. As noted in the introduction, self-disclosure can vary according to the *breadth* and *depth* of information disclosed. The two techniques studied in the present paper outline alternative methods to measure both aspects of disclosure. In the case of an 'I prefer not to say' option, the breadth of disclosure can be studied by examining the number of items people choose to disclose or non-disclose to. In Study 1, we reported that this use of 'I prefer not to say' is sensitive to the priming of privacy concerns, and in Study 2, that it is sensitive to the experimental manipulation of privacy. A proxy measure of the depth of disclosure would be the inclusion of items of varying sensitivity. In Study 1, we reported that non-disclosure was higher to sensitive items compared to non-sensitive items, although both forms were responsive to the priming of privacy.

However, there are two potential issues with the use of 'I prefer not to say' response options to measure disclosure online. The first is that, in some cases, choosing the 'I prefer not to say' option may imply a response that is socially undesirable. For instance, responding 'I prefer not to say' to a question such as, "have you ever been arrested for a violent act?" may in itself imply that the respondent has something to hide. For this reason, we would caution that, as a measure of self-disclosure, this technique should only be used for items that do not imply a response through the use of 'I prefer not to say'. A second issue is that variance in the use of 'I prefer not to say' was low, with relatively few people choosing the option. Both sensitive and non-sensitive items were responsive to the privacy priming in Study 1, suggesting that non-disclosure may be influenced by the formation of a disclosure/non-disclosure response set. For this reason, the use of "I prefer not to say" may be best treated as a dichotomous variable, rather than as continuous, although further research is needed to confirm this.

The second technique we studied was the provision of an opportunity to vary the granularity of information disclosed through blurring. This technique could be used for a variety of questions where a range response is possible (e.g. income, units of alcohol and sexual partners). In Study 2, we found evidence that the opportunity to 'blur' information was taken specifically by males in the high privacy condition, when compared to females in both conditions, and males in the low privacy condition. This suggests that our technique for the measurement of blurring could potentially act as a valuable measure of the depth of disclosure. While the decision whether or not to disclose reflects the breadth of disclosure, the granularity of the information disclosed reflects the depth of disclosure. In the high privacy condition, males may have made the strategic decision to disclose in light of the privacy assurances they received, but relied on blurring to increase the ambiguity

of information provided as an alternative privacy protection technique. The results of Study 2 highlight the importance of studying both the breadth and depth of disclosure in that the group with the highest breadth of disclosure (males in the high privacy condition) was also the group who showed the lowest depth in terms of blurring.

While it is well established that males and females have different levels of self-disclosure (Dindia & Allen, 1992), the effect is usually that women disclose somewhat more than males (see Dindia & Allen, 1992, for a meta-analysis). However, it is not clear if these differences extend to computer-mediated disclosure and in particular privacy concerns, although Garbarino and Strahilevitz (2004) found that women, compared to males, tended to view the risks of online shopping (including loss of privacy) as more problematic. In the present study, women's lower levels of disclosure might reflect the nature of the questions (e.g. sexual behavior), or increased concern with privacy issues. In Study 2, women were less likely to disclose their income compared to males. This finding ran counter to our expectation that income would be more sensitive for males compared to females. There are a number of possible explanations for this. The first is that, given the sample (part-time adult students), income (or lack of) might be a more sensitive subject for females than males. A second explanation could be that, for males, discussions of income are more commonplace, reducing the sensitivity of the question. Finally, given continuing inequality in remuneration, the topic of income may simply be more sensitive for females compared to males. Certainly, future research should be conducted to examine if this effect generalizes outside of the specific sample in the present study. Future research could also investigate gender differences in responses to privacy interventions alongside the use of different disclosure strategies. While it is reasonably well established that many people do not read privacy policies (Berendt, Gunther, & Spiekerman, 2005; Milne & Culnan, 2004), the results of the present research would suggest that adding a gender dimension, alongside behavioral measures of disclosure, would be potentially valuable.

Self-disclosure is a critical issue in the design of trustworthy, privacy protecting web sites and services. It is also of particular interest to researchers who wish to study the impact of systems design on human–computer interaction. The techniques outlined in this paper provide researchers with ways to easily measure people's willingness to disclose information in such environments behaviorally, rather than their reported or intended self-disclosure. The results of the analyses of non-disclosure and blurring in response to the income question by males and females suggest that both techniques are valuable additions to the range of methods available for the study of people's privacy-related behaviors when interacting with computer systems.

Acknowledgements

References

Altman, I., & Taylor, D. (1973). *Social penetration: The development of interpersonal relationships*. New York: Holt, Rinehart and Winston.

Antaki, C., Barnes, R., & Leudar, I. (2005). Self-disclosure as a situated interactional practice. *British Journal of Social Psychology, 44*, 181–200.

Bargh, J. A., McKenna, K. Y. A., & Fitzsimons, G. M. (2002). Can you see the real me? Activation and expression of the true self on the Internet. *Journal of Social Issues, 58*, 33–48.

Berendt, B., Gunther, O., & Spiekerman, S. (2005). Privacy in E-commerce: Stated behaviour versus actual behaviour. *Communications of the ACM, 48*, 101–106.

Buchanan, T., Paine, C. B., Joinson, A. N., & 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.

Buchanan, T., Joinson, A. N., & Ali, T. (2002). Development of a behavioural measure of self-disclosure for use in online research. In *Paper presented at German Online Research '02, Hohenheim, Germany*.

Chellappa, R. K., & Sin, R. G. (2005). Personalisation versus privacy: An empirical examination of the online consumer's dilemma. *Information Technology and Management, 6*, 181–202.

Culnan, M. J. (1999). *Georgetown Internet Privacy Policy Study: Privacy Online in 1999: A report to the Federal Trade Commission*. Washington DC: Georgetown University.

Dindia, K., & Allen, M. (1992). Sex differences in disclosure: A meta-analysis. *Psychological Bulletin, 112*, 106–124.

Duckham, M., & Kulik, L. (2005). A formal model of obfuscation and negotiation for location privacy. In *3rd International conference on pervasive computing: Pervasive '05*. Munich, Germany: Springer.

Garbarino, E., & Strahilevitz, M. (2004). Gender difference in the perceived risk of buying online and the effects of receiving a site recommendation. *Journal of Business Research, 57*(July), 768–775.

Galegher, J., Sproull, L., & Kiesler, S. (1998). Legitimacy, authority and community in electronic support groups. *Written Communication, 15*, 493–530.

Gruteser, M., & Grunwald, D. (2003). Anonymous usage of location-based services through spatial and temporal cloaking. In *First international conference on mobile systems, applications, and services*.

Holbrook, A. L., Green, M. C., & Krosnick, J. A. (2003). Telephone vs. face-to-face interviewing of national probability samples with long questionnaires: Comparisons of respondent satisficing and social desirability response bias. *Public Opinion Quarterly, 67*, 79–125.

Joinson, A. N. (1999). Social desirability, anonymity and Internet-based questionnaires. *Behavior Research Methods, Instruments and Computers, 31*(3), 433–438.

Joinson, A. N. (2001a). 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. (2001b). Knowing me, knowing you: Reciprocal self-disclosure on the Internet. *Cyberpsychology and Behavior, 4*, 587–591.

Joinson, A. N., & Paine, C. B. (2007). Self-disclosure, privacy and the Internet. In A. N. Joinson, K. Y. A. McKenna, T. Postmes, & U-D. Reips (Eds.), *Oxford handbook of Internet psychology* (pp. 237–252). Oxford University Press.

Joinson, A. N., Woodley, A., & Reips, U.-D. (2007). Personalization, authentication and self-disclosure in self-administered Internet surveys. *Computers in Human Behavior, 23*(1), 275–285.

Joinson, A. N., & Reips, U.-D. (2007). Personalized salutation, power of sender and response rates to Web-based survey. *Computers in Human Behavior, 23*(3), 1372–1383.

Jourard, S. M. (1971). *Self-disclosure: An experimental analysis of the transparent self*. New York: Krieger.

Jourard, S. M., & Lasakow, P. (1958). Some factors in self-disclosure. *Journal of Abnormal and Social Psychology, 56*(1), 91–98.

Knapp, H., & Kirk, S. A. (2003). Using pencil and paper, Internet and touch-tone phones for self-administered surveys: Does methodology matter? *Computers in Human Behavior, 19*, 117–134.

Laurenceau, J. P., Barrett, L. F., & Pietromonaco, P. R. (1998). Intimacy as an interpersonal process: The importance of self-disclosure, partner disclosure, and perceived partner responsiveness in interpersonal exchanges. *Journal of Personality and Social Psychology, 74*, 1238–1251.

Metzger, M. J. (2004). Privacy, trust and disclosure: Exploring barriers to electronic commerce. *Journal of Computer-Mediated Communication, 9*(4). <http://jcmc.indiana.edu/vol9/issue4/metzger.html>. Accessed 20.06.05.

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*(3), 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.

Nickel, J., & Schaumburg, H. (2004). Electronic privacy, trust and self-disclosure in E-recruitment. In *Proceedings of CHI 2004, April 24–29, Vienna*.

Nyshadham, E. A. (2000). Privacy policies of air travel web sites: A survey and analysis. *Journal of Air Transport Management, 6*, 143–152.

Parks, M. R., & Floyd, K. (1996). Making Friends in Cyberspace. *Journal of Communications, 46*, 80–97.

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. In *Proceedings of the third ACM conference on electronic commerce, Association for Computing Machinery (ACM EC'01), Tampa, Florida, US*, pp. 38–47.

Suler, J. R. (2004). The online disinhibition effect. *Cyber Psychology and Behavior, 7*, 321–326.

Tourangeau, R. (2004). Survey research and societal change. *Annual Review of Psychology, 55*, 775–801.

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.