

# Experimental Studies of Disclosure Risk, Disclosure Harm, Topic Sensitivity, and Survey Participation

*Mick P. Couper, Eleanor Singer, Frederick G. Conrad, and Robert M. Groves*<sup>1</sup>

This article extends earlier work (Couper et al. 2008) that explores how survey topic and risk of identity and attribute disclosure, along with mention of possible harms resulting from such disclosure, affect survey participation. The first study uses web-based vignettes to examine respondents' expressed willingness to participate in the hypothetical surveys described, whereas the second study uses a mail survey to examine actual participation. Results are consistent with the earlier experiments. In general, we find that under normal survey conditions, specific information about the risk of identity or attribute disclosure influences neither respondents' expressed willingness to participate in a hypothetical survey nor their actual participation in a real survey. However, when the possible harm resulting from disclosure is made explicit, the effect on response becomes significant. In addition, sensitivity of the survey topic is a consistent and strong predictor of both expressed willingness to participate and actual participation.

*Key words:* Disclosure; informed consent; survey participation; confidentiality; topic sensitivity.

## 1. Introduction

The research reported in the present article is part of a research program aiming to estimate the risk (probability) of statistical disclosure in publicly available data files and to develop procedures for communicating this risk to potential respondents. "Statistical disclosure" refers to the ability to deduce an individual's identity despite the absence of personal identifiers such as name and address on the data file, through a process of matching an individual's de-identified record against another record containing (some of) the same characteristics as the original file in addition to the person's name and address. Although successful matches have been reported (e.g., Paass 1988; Malin and Sweeney 2000; Winkler 1997), little is known about the likelihood of such disclosures from publicly available data files.

The ultimate aim of our part of this research program was to use empirically determined estimates of disclosure risk (i.e., a specific probability that respondents' names will be associated with their answers), which we expected to be very low, to craft informed

<sup>1</sup> Survey Research Center, Institute for Social Research, University of Michigan, P.O. Box 1248, Ann Arbor, MI 48106-1248, U.S.A. Emails: mcouper@isr.umich.edu, esinger@isr.umich.edu, fconrad@isr.umich.edu, and Bgroves@isr.umich.edu

**Acknowledgments:** We thank NICHD (Grant #P01 HD045753-01) for support. We thank Reg Baker for help with deploying the Web survey, and Rachel Orłowski and Ruth Philippou for deploying the mail survey. We also thank John Van Hoeyk for indispensable help with analyses and Rachel Orłowski and Catherine Militello for coding the open-ended responses.

consent statements that would meet ethical requirements without unnecessarily arousing respondents' concerns about the confidentiality of their personal information, since much earlier research has indicated that concerns about confidentiality are capable of substantially reducing survey participation (for a summary, see Couper et al. 2008).

We have so far carried out five experimental studies as part of this research program. The first two were carried out in the lab and were reported in Conrad et al. (2006) and summarized in Couper et al. (2008). Briefly, they suggested that only extreme disclosure risk (i.e., virtual certainty of disclosure) affected (negatively) willingness to participate in a hypothetical study, and that describing a study as dealing with a topic most would consider sensitive increased the impact of the risk manipulation.

The third study, a web-based vignette experiment (Couper et al. 2008), was designed to investigate the effects of topic sensitivity and risk of identity and attribute disclosure on respondents' expressed willingness to participate in the survey described. We found that under conditions resembling those of real surveys, objective risk information did not affect expressed willingness to participate. On the other hand, topic sensitivity did have such effects, as did general attitudes toward privacy and survey organizations as well as subjective perceptions of risk, harm, and benefits.

The present article describes the fourth and fifth studies of this research program. The fourth is a web experiment to test some implications of earlier experiments, and the fifth validates the laboratory and web experiments in a real-world setting, using a mail survey.

## **2. Web Experiment: The Effect of Disclosure Risk and Harm**

Given the modest effects of the risk of disclosure manipulation on willingness to participate, we speculated that it might be the risk of harm resulting from disclosure, rather than the risk of disclosure itself, that affected participation. We thus designed a second web experiment that made explicit not only the risk of disclosure but also the harm that might follow on disclosure of respondents' identity along with their answers. We also attempted to make salient privacy concerns in one experimental condition and to avoid doing so in another, because the first web experiment had found greater sensitivity to risk among respondents most concerned about privacy. Finally, we varied the size of the incentive respondents were offered, and the mode of the survey request. Other elements of the original web experiment were held constant, but the sample size was increased.

We hypothesized the following effects:

1. Making harm of disclosure explicit will result in a significant effect of disclosure risk on stated willingness to participate;
2. Making privacy concerns salient will reduce willingness to participate and increase the impact of the risk manipulation;
3. Larger incentives will result in greater willingness to participate, but in line with earlier research there will be no interaction between the size of the incentive and the size of the disclosure risk the respondent is willing to assume;
4. Because of the greater privacy it affords, a mail survey request will result in greater willingness to participate than a face-to-face survey request;
5. Sensitive topics will reduce willingness to participate and increase the effect of disclosure risk on participation, relative to nonsensitive topics.

## 2.1. Methods

The experiment was designed as a web survey, which was administered by Market Strategies, Inc., on a volunteer sample drawn from Survey Sampling International's Internet panel. We received 6,400 completed questionnaires out of a total of 217,542 invitations sent, for a "response rate" of 2.9 percent. Invitees were told it was a study on survey participation, and that they would see descriptions of different types of surveys and be asked whether or not they would be likely to take part.

Respondents were asked at the end of the survey whether they had noticed any differences among the vignettes, and if so, what differences they noticed. Of those asked, 4,603 (or 72%) reported noticing differences among the vignettes, and we restrict our analyses to these respondents.<sup>2</sup> Of these, 54% were female, 19% had a high school education or less and 41% reported being college graduates, 77% were nonHispanic whites, 23% were under 30 years old and 42% older than 50. This is not a probability sample. Our focus, to use Kish's (1987) terms, is on randomization rather than representation. We view this as an experiment with a large and diverse group of volunteer subjects.

Each questionnaire included a set of eight fictional survey invitations, or vignettes (described below); a question about willingness to participate in each of these fictional surveys; and other questions designed to explore perceptions of risk and benefit as well as concerns about privacy and attitudes toward surveys, which are identical to those in the earlier web study (Couper et al. 2008). The entire questionnaire took an average of 16 minutes to complete.

The vignettes experimentally varied four factors: the survey topic; the description of the risk (probability) of disclosure together with a description of the harm such disclosure might cause; the size of the incentive for participation (\$10 or \$50); and the mode (face-to-face or mail). In addition, half the sample had privacy made salient at the outset by asking a few questions related to privacy concerns ("privacy prime") while the other half was given the same number of questions about a neutral topic, computer use ("neutral prime"). For a given vignette, disclosure risk was set at one of four levels: No mention; no chance of disclosure; one in a million; and one in ten. Two levels of topic sensitivity, each with two specific topics, were varied across the vignettes. The high-sensitivity topics were sexual behavior and personal finances; the low-sensitivity topics were leisure activities and work. Each vignette included a confidentiality assurance: "The information you provide is confidential." Each vignette also mentioned the fictional study's sponsor (the National Institutes of Health), a benefit statement tailored to the topic (see the sample vignette below for an example), and the estimated survey length (20 minutes); these features were kept constant across all 32 vignettes resulting from the complete crossing of Topic (4) × Risk (4) × Incentive (2). The privacy prime and mode were between-subject manipulations.

As noted, each subject was exposed to a subset of 8 of these vignettes, with each set containing risk descriptions at all four levels, one each for a sensitive and a nonsensitive

<sup>2</sup> Including these cases in the analysis does not substantially alter any of the results or conclusions.

topic. The sets were randomly assigned to subjects after they had agreed to participate in the Web survey, and the order in which the vignettes were administered was also randomized within subjects. Analyses are presented for all eight vignettes and for the first vignette only; we regard the latter as a more rigorous test of effects because it prevents respondents from using the other vignettes as reference points. An example vignette is shown below:

“Imagine that you have cheated on your partner during the past year by having sex with another person and that a professional survey interviewer visits your home and says the following:

*My name is Mary Jones and I work for the University of Michigan Survey Research Center. We would like you to take part in a survey on sexual behavior and sexually transmitted diseases, sponsored by the National Institutes of Health. The information you provide will help shape government policy on sexually transmitted diseases.*

*The information you provide is confidential. Based on experience, we think there is a one in 10 chance that someone will connect your name with your answers.*

*The interview will take 20 minutes, and you will receive \$10 as a token of the researcher’s appreciation.*

If your partner connects your name with your answers on the survey, this might result in the break-up of your relationship.”

Each hypothetical harm was tailored to the specific topic of the vignette on the basis of discussions among the investigators. In the vignette describing a survey about sexual behavior, above, the harm from disclosure was described as the potential break-up of the relationship; in the vignette about financial assets, the harm was described as potential discovery of tax cheating by the IRS and a resulting fine; in the vignette about a leisure time survey, the harm was described as surprise on the part of friends who might discover the participants’ amount of TV watching; in the vignette about a survey of working conditions, the harm was described as receipt of additional junk mail. We tried to vary the severity of the harms to parallel the sensitivity of the vignettes, and to identify potential harms applicable to a broad range of respondents.

Immediately following each vignette, respondents were asked how likely they would be to participate in the survey described, on an 11-point scale from 0, not at all likely, to 10, very likely. For the first vignette only, this question was followed by an open-ended question, Why or Why not, and then by questions about perceived risk, perceived harm, perceived personal and social benefits, and the perceived risk/benefit ratio. These measures are described in detail in Couper et al. (2008).

In addition to the vignettes and related questions, the web questionnaire included questions about privacy, attitudes toward surveys, and trust. These questions, asked of everyone, are also described in Couper et al. (2008).

## 2.2. Results

Table 1 presents models addressing hypotheses 1–5. A control variable is added to the model using all eight vignettes to account for learning effects (vignette order). The first model is estimated using OLS regression, while the second is estimated using SAS PROC

Table 1. The Effect of Key Experimental Manipulations on Willingness to Participate (Web Experiment)

	First vignette		All vignettes	
	Coefficient	(Std. err.)	Coefficient	(Std. err.)
Intercept	7.993***	(0.186)	7.668***	(0.088)
Risk				
One in ten	-0.287	(0.236)	-0.485***	(0.067)
One in a million	0.162	(0.237)	-0.0083	(0.067)
No chance	0.114	(0.252)	-0.036	(0.067)
No mention	-	-	-	-
Sensitivity (1=high)	-2.724***	(0.199)	-3.009***	(0.055)
Prime (1=privacy)	-0.440***	(0.199)	-0.429***	(0.086)
Incentive (1=high)	0.285**	(0.098)	0.270***	(0.072)
Mode (1=FTF)	-0.184	(0.098)	-0.324***	(0.072)
Risk × sensitivity				
1/10, high	-0.276	(0.279)	-0.221**	(0.078)
1/million, high	0.433	(0.277)	0.076	(0.078)
No chance, high	-0.056	(0.281)	-0.090	(0.078)
Risk × privacy prime				
1/10, privacy	-0.278	(0.279)	-0.070	(0.078)
1/million, privacy	-0.247	(0.278)	-0.083	(0.078)
No chance, privacy	-0.078	(0.282)	-0.028	(0.078)
Vignette number			-0.036***	(0.006)
Observations	4,587		36,584	
Model adjusted R <sup>2</sup>	0.153		0.174	

MIXED and IVEware (Raghunathan et al. 2001, 2005) to control for the repeated measures within person. As already noted, all mentions of risk in this study also included mention of specific harms associated with that risk.

First, the main effect of disclosure risk is significant ( $F[3, 4, 573]=6.26, p < .0001$  for the first vignette and  $F[3, 31, 966]=114.41, p < .0001$  for all eight vignettes). The greater the risk, the lower respondents' WTP (see Table 1). Similar results are obtained for models without the interaction terms (not shown). These results can be contrasted with our earlier study (Couper et al. 2008) where the vignettes did not mention harm. In that study, only the analyses based on responses to all 8 vignettes yielded a significant effect of the risk manipulation, presumably because respondents were able to use the various vignettes as reference or comparison points.<sup>3</sup> In both models in Table 1, the highest risk (1 in 10), now coupled with mention of the likely harm, yields significantly lower levels of stated willingness to participate (WTP) than do all other disclosure risk categories. Thus, our first hypothesis is supported.

Turning to our second hypothesis, we find significant negative effects of the privacy prime—relative to the neutral (computer use) prime—in both models. That is, those for whom privacy was made salient have significantly lower levels of WTP than those given the neutral prime. In analyses not shown, we also found that the privacy prime had

<sup>3</sup> For an early experiment demonstrating the development of “subjective norms” over a series of judgments when objective standards of comparison are lacking, see Sherif (1958).

significant effects on perceived risk and perceived harm, and on general attitudes toward privacy, all measured following the vignettes. This suggests that the privacy prime was effective in raising general privacy concerns. However, we found no significant effect of the interaction of the prime with the risk manipulation on WTP ( $F[3, 4, 573]=5.03$ ,  $p=0.71$  for the first vignette and  $F[3, 31, 966]=0.48$ ,  $p=0.70$  for all vignettes). That is, contrary to our expectation, those who received the privacy prime were no more susceptible to the risk manipulation than those who got the neutral prime.

Our third hypothesis relates to the effect of the incentive. We compared a large (\$50) with a small (\$10) incentive. As can be seen from Table 1, there is a significant positive effect of the incentive in both models—the larger incentive is associated with higher levels of stated WTP. In separate analyses (Couper and Singer 2009) we found no evidence that respondents were willing to trade higher risks for larger incentives; the interaction between risk and incentive size is not significant. The overall interaction between risk and incentive is not significant in either model ( $F[3, 4, 579]=0.73$ ,  $p=0.54$  for the first vignette;  $F[3, 31, 969]=1.11$ ,  $p=0.34$  for all eight vignettes).

The fourth hypothesis, relating to the effect of mode, was partially supported by the results in Table 1. In both models, those in the face-to-face condition expressed lower levels of WTP than those in the mail condition, although this difference fails to reach conventional levels of significance ( $p=.09$ ) in the first vignette model.

The final hypothesis relates to the effect of topic sensitivity on WTP. As before, we found a strong main effect for sensitivity of the topic, with significantly lower WTP for the more sensitive topics (sexual behavior and personal finances) than for the less sensitive ones (leisure activities and work) (see Table 1). However, our focus is on the interaction between sensitivity and risk. We hypothesized that the effect of the risk manipulation would be stronger for more sensitive topics than for less sensitive ones. Overall, the interaction term is not significant in the first vignette model ( $F[3, 4, 573]=2.34$ ,  $p=.071$ ) but reaches conventional levels of significance in the model with all eight vignettes ( $F[3, 31, 966]=5.35$ ,  $p=.0011$ ). However, an examination of individual contrasts in the first vignette model using the Tukey-Kramer adjustment for multiple contrasts shows the difference in WTP between the 1/ten and 1/million groups to be statistically significant for the more sensitive topics ( $t=5.98$ ,  $p < .001$ ) but not for the less sensitive topics ( $t=2.40$ ,  $p=0.24$ ).

The interaction is plotted in Figure 1, using the least squares means from the first vignette model. The strong main effect of topic sensitivity can be seen with the high sensitivity topics (solid line) significantly lower on WTP than the low sensitivity topics (dashed line), across all levels of risk. The effect of the risk manipulation is also seen, with the 1/ten group differing from the other three groups. Finally, the slopes between the 1/10 and 1/million groups are significantly different, with a larger effect (steeper slope) for the more sensitive topics. These results are mirrored in the findings for the model based on all eight vignettes. We thus find support for the final hypothesis that the effect of the risk manipulation on WTP depends on the sensitivity of the topic.

We also added a set of general attitudes and demographic controls to the models in Table 1, to see if these affect the key relationships explored above. The main effect of the topic sensitivity indicator remains significant in both models ( $F[1, 4, 477]=305.2$ ,  $p < .0001$  for the first vignette, and  $F[1, 30, 755]=12,040.9$ ,  $p < .0001$  for all eight

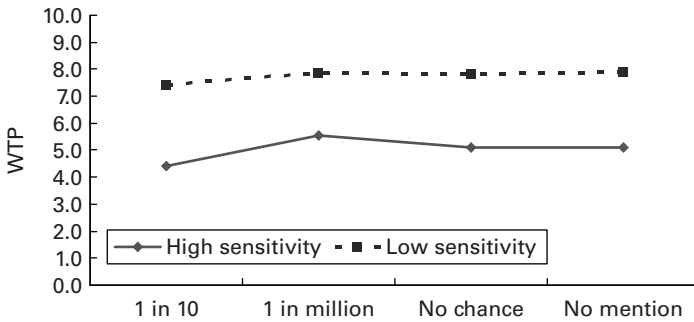


Fig. 1. Risk\*Sensitivity Interaction, First Vignette (Web Experiment)

vignettes), as does the main effect of the risk manipulation ( $F[3, 4, 456]=5.17, p=.0015$  for the first vignette, and  $F[3, 30, 755]=112.19, p<.0001$  for all eight vignettes). The risk  $\times$  sensitivity interaction approaches conventional significance levels for the first vignette ( $F[3, 4, 456]=2.54, p=.054$ ) and is significant in the model with all eight vignettes ( $F[3, 30, 755]=5.38, p=.0011$ ). The direction of the effect mirrors that for the first vignette shown in Figure 1. Similarly, the main effects of mode, incentive, and the primacy prime remain statistically significant in the presence of the general attitude measures and demographic controls. This provides further support for the hypotheses outlined earlier.

As already noted, we elicited open-ended comments in response to the question, “Why would you/would you not be willing to participate in the survey described?” By far the most frequent reasons for not participating in the survey were coded as privacy-related (e.g., “don’t like intrusions”; “too much personal information”; “don’t want strangers in home”; “want/don’t trust confidentiality assurance”)—62.8% of reasons given fell into this category, compared with 47.4% who had cited privacy-related concerns in the earlier experiment that made no mention of harm (Couper et al. 2008). The next largest category of objections related to aspects of the survey (e.g., “takes too long”; “need to know more about questions”), which were mentioned by 14.8% of the sample.

In summary, the results of the second web experiment, compared with those described earlier (Couper et al. 2008), suggest that when there is explicit mention of the likely harm associated with the risk, respondents’ willingness to participate in hypothetical surveys can be affected by the stated risk of disclosure even when they are exposed to only a single vignette. Except when harm is explicitly mentioned, respondents appear to be indifferent to variations in disclosure risk that may be mentioned as part of a survey request. Note, however, that since we tested only three specific risk variations, the generalizability of this conclusion may be limited.

The web experiment was conducted among a generally cooperative group of respondents—those who had agreed to join an opt-in panel and made the decision to complete the web survey containing our experiment—and measured only willingness to participate in a hypothetical survey rather than actual participation. To remedy these limitations, we carried out a fifth experiment, explicitly designed to test our hypotheses in a real-life setting.

### 3. Mail Experiment: Validating the Results of the Earlier Studies

The validation study was carried out as a mail survey, partly for reasons of cost, but partly because we did not want interviewers to contaminate the experimental manipulations. Based on the results of the laboratory and web studies described above, we formulated the following hypotheses for the mail survey:

1. Sensitive topics (sex, money) will produce lower response rates than non-sensitive topics (leisure, work).
2. Mention of possible harm from disclosure will significantly decrease response.
3. The disclosure manipulation will interact with topic and mention of harm, so that when the topic is sensitive, the risk-harm manipulation will reduce response rates; whereas when the topic is not sensitive, the risk-harm manipulation will have no effect.

On the basis of the earlier experiments, we expected no main effect of disclosure risk on response.

#### 3.1. Design

The mail survey varied the three elements identified as significant in the earlier studies: topic, risk of disclosure, and possible harm from disclosure. Crossing the risk and harm dimensions used previously yields the following five conditions:<sup>4</sup>

- High risk (1 in 10), mention of harm
- High risk (1 in 10), no mention of harm
- Low risk (1 in 1,000,000), mention of harm
- Low risk (1 in 1,000,000), no mention of harm
- No mention of risk or harm

Crossing the five risk-harm conditions with the four topics previously tested (sexual behavior, finances, work, and leisure activities) yields a  $5 \times 4 = 20$  condition experimental design.

The experimental manipulations were created by means of 20 different cover letters mailed to respondents, which varied the description of the topic, disclosure risk, and harm likely to be associated with the survey. For example, the letter below was sent to sample members in the Sex  $\times$  High risk  $\times$  Harm condition:

*June 20, 2008*

*Dear First Name Middle Name Last Name*

*We are writing to invite your participation in a national survey on how sexual behavior and sexually transmitted diseases are related to health and well-being. This survey is being conducted by the Survey Methodology Program, Institute for Social Research. A questionnaire and a stamped return envelope are enclosed with this letter, along with*

<sup>4</sup>The zero-risk condition used in the vignette studies was omitted because it is impossible to guarantee in an actual survey.



*\$2 as a token of our appreciation. Filling out the questionnaire should take no more than about 15 minutes.*

*Although you may find some of the questions sensitive, the information we request is vitally important. The responses to the survey will help to inform policy on these important topics.*

*The information you provide is confidential. Based on experience, we think there is a one in ten chance that someone will connect your name with your answers. If this were to occur, here are some of the things that might happen as a result:*

- *A spouse or partner finds out that the respondent has been cheating on them;*
- *An insurer discovers that the respondent has a sexually transmitted disease and raises the cost of health care coverage;*
- *Friends or family discover that the respondent engages in risky sexual behavior, and react with disapproval.*

*The sample for the survey has been selected at random from a listing of all U.S. households. In order for the results to be of scientific value, it is very important that everyone selected for the study respond. Because you cannot be replaced with anyone else, we hope that we can count on you. If you have questions, please call us toll-free at 1-800-0000 or email us at [SMPsurvey@isr.umich.edu](mailto:SMPsurvey@isr.umich.edu). Should you have questions regarding your rights as a participant in research, please contact the Institutional Review Board – Behavioral Sciences, 540 East Liberty Street, Suite 202, Ann Arbor, MI 48104-2210; Phone: 734-936-0933; e-mail: [irbhsbs@umich.edu](mailto:irbhsbs@umich.edu).*

Note that in the mail survey, as distinct from the web vignettes, several possible outcomes varying in severity are linked to each topic in order to encourage more respondents to identify with the possible harmful outcomes of disclosure and to avoid the possibility that the findings are limited to a single specific harm associated with each topic.<sup>5</sup>

The questionnaire consisted of a twelve-page booklet containing both unique and common elements. The first section contained unique, topic-related questions that varied from one booklet to another, depending on the particular topic to which a respondent was assigned; the remaining sections, which asked questions about general health and well-being, concerns about privacy and confidentiality, attitudes toward surveys, and demographic characteristics, were common across all conditions of the experiment.

<sup>5</sup> We have measured harm in a variety of ways. In an earlier study, Singer (2003) measured perceived harm by asking respondents to a telephone survey using hypothetical vignettes what they thought the likelihood of disclosure of their name and address, along with their answers, would be to various groups (family, friends, employers, and law enforcement agencies); she then followed this with questions about how much they would mind if each of these groups discovered this information. In the Web experiment, we gave respondents an example of a single harm that might follow on disclosure. In the mail experiment, we gave respondents either two (for nonsensitive vignettes) or three (for sensitive vignettes) examples of possible harms. Each of these definitions of harm resulted in a large and significant reduction in either willingness to participate or actual participation.

### 3.2. Sample and Data Collection

A national sample of 8,000 names was purchased from Genesys Sampling Systems, selected from those for whom at least a surname and an address were available. Sampled persons were assigned to the twenty treatment cells systematically, assuring a nearly uniform distribution by census region, age group, and gender (where this information was available); sampled persons without values for age and gender were then randomly distributed by region. This resulted in 2,000 sampled persons per topic and 1,600 per risk-harm condition. Since the main dependent variable was the response rate, these numbers provided ample analytic power.

The cover letter and questionnaires were mailed on June 20, 2008, along with a \$2 cash incentive, followed by a reminder postcard on July 15. A reminder letter with a replacement questionnaire was mailed on September 2, 2008, and a nonresponse follow-up letter including a postcard requesting precoded reasons for nonresponse was mailed on October 15 to all those who had not responded and for whom we had not received an “undeliverable” notification. A final debriefing letter, which disclosed the experimental nature of the study and assured sampled persons of the confidentiality of their answers, was mailed on December 1, 2008 to all those for whom we had a deliverable address. This is comparable to the procedure that would be employed in any laboratory study involving deception.

### 3.3. Results

Overall, 1,831 completed questionnaires and 72 sufficient partials (i.e., having completed at least Section A) were returned, for a response rate of 23.8% and a return rate of 28.3% (restricted to the 6,723 sample addresses exposed to the manipulation—i.e., excluding all Postmaster Returns, undelivered questionnaires, and ineligible respondents).

The return rate (completed + sufficient partials divided by those exposed to the manipulation) for each of the cells in the experimental design is shown in Table 2. Similar results are obtained when excluding partials or retaining all 8,000 sample cases in the denominator. We used logistic regression to test the significance of the main and interaction effects we hypothesized. We first fit a model with the main effects of topic (4 categories) and risk-harm (5 categories) and tested for the contrasts implied by the hypotheses. We then added the topic  $\times$  risk-harm interaction in a second model.

We find support for the first hypothesis: return rates are lower for sensitive topics (sex and money) than for nonsensitive topics. The overall test of topic is statistically significant (Wald  $\chi^2[3]=255.44$ ,  $p < .0001$ ), as is the contrast between sensitive and nonsensitive topics (Wald  $\chi^2[1]=223.48$ ,  $p < .0001$ ). Turning to the second hypothesis, we find a main effect of the mention of harm (Wald  $\chi^2[1]=13.8$ ,  $p = .0002$ ), in the expected direction. As we expected, we found no main effect of disclosure risk on return (Wald  $\chi^2[2]=3.29$ ,  $p = .19$ ). These results are consistent with the findings from the earlier web surveys. We also combined the risk-harm manipulation into a single five-category variable, which does reach statistical significance (Wald  $\chi^2[4]=23.33$ ,  $p = .0001$ ), but as can be seen from Table 2 and the separate tests above, this is largely driven by the mention of harm associated with the risk, rather than with the level of disclosure risk mentioned.

Table 2. Return Rates by Topic and Risk-Harm Manipulation (Mail Experiment)

Risk-harm manipulation	Topic				Total
	Sensitive		Nonsensitive		
	Sex	Money	Leisure	Work	
High risk, harm	16.4% (335)	15.7% (332)	34.8% (339)	28.8% (333)	24.0% (1,339)
High risk, no harm	25.8% (329)	20.7% (348)	44.7% (333)	36.6% (342)	31.9% (1,352)
Low risk, harm	23.5% (323)	15.5% (329)	46.1% (321)	29.9% (341)	28.7% (1,314)
Low risk, no harm	23.5% (328)	18.9% (334)	42.9% (352)	33.3% (345)	29.9% (1,359)
No mention of risk or harm	22.3% (337)	20.1% (338)	36.7% (349)	35.2% (335)	28.6% (1,359)
Total	22.3% (1,652)	18.2% (1,681)	41.0% (1,694)	32.8% (1,696)	28.6% (6,723)

The final hypothesis deals with an interaction between the combined risk-harm variable and topic. The test of the interaction does not reach statistical significance, either when examining the four topics separately (Wald  $\chi^2[12]=12.47, p=.409$ ) or when collapsing the topics into high versus low sensitivity (Wald  $\chi^2[4]=1.29, p=.86$ ). We thus find no support for the final hypothesis that the effect of the risk-harm manipulation would vary with topic sensitivity.

The return rates for high/low sensitivity by risk-harm are shown in Figure 2.

The effect of topic sensitivity on mail survey return is clear from this figure. As we have already noted, the main effect of topic sensitivity is significant. The effect of the mention of the possible harms associated with disclosure is also apparent from the figure: return rate rates are higher when there is no mention of harm. In other words, it is the mention of harm that affects willingness to participate, rather than the mention of the disclosure risk.




Fig. 2. Risk-Harm\*Sensitivity Interaction (Mail Experiment)

In addition to examining return rates, we created an index of item missing data among respondents for the first section of the questionnaire, containing the questions germane to the stated topic. Given the skewed nature of these distributions, we ran a Poisson regression of missing data rates on topic and risk-harm. The results of this model (not shown) parallel the findings for return rates: there is a significant main effect (Wald  $\chi^2[3]=23.75$ ,  $p < .0001$ ) of topic on item missing data rates (with more missing data for the more sensitive topics), but no effect for the risk-harm manipulation (Wald  $\chi^2[4]=0.45$ ,  $p=.97$ ), and no significant interaction (Wald  $\chi^2[12]=5.11$ ,  $p=.95$ ). In other words, given that a questionnaire was returned, we find no differences in missing data rates within topic across the different risk-harm manipulations.

The final analysis examined the responses to the nonresponse letter and postcard. Overall, 376 postcards were returned with reasons for nonresponse, out of 5,284 sent, for a response rate of 7.1%. The postcard response rate did not differ by topic or risk-harm manipulation. Among those who returned the postcards, there were no differences in reasons given for nonresponse by risk-harm. However, topic did matter: 33% of those who received questionnaires on the more sensitive topics (sex and money) checked "Topic too sensitive" as a reason for not participating, compared to 7% of who received questionnaires on the less sensitive topics (work and leisure) (Wald  $\chi^2[1]=29.27$ ,  $p < .0001$ ). Similarly, significantly more of those in the more sensitive topic condition checked "I was concerned someone might see my answers" (13.5% versus 5.2%; Wald  $\chi^2[1]=6.35$ ,  $p=.012$ ) and "I was afraid my answers might be used against me" (12.1% versus 5.9%).

#### **4. Discussion and Conclusions**

We began this research with what, in retrospect, seems like a very naïve hypothesis: Giving potential respondents accurate information about the likelihood that their information would be revealed to unauthorized persons, primarily through statistical disclosure, would not only serve the requirements of informed consent but would also serve to reassure respondents, since we expected the likelihood of disclosure to be vanishingly small.

Our speculations turned out to be naïve for two reasons. First, the task of estimating the likelihood of statistical disclosure is extremely complicated and not easily generalized from one survey to another. But second, even if it were possible to provide respondents with such information, our research has indicated that they are generally impervious to it. Although we were initially skeptical about this conclusion, repeated experiments in a variety of modes have persuaded us that under normal survey conditions, like those in our mail survey, specific information about the risk that respondents' names and addresses will be linked to their answers influences neither respondents' expressed willingness to participate in a hypothetical survey nor their actual participation in a mail survey. Nor, finally, does it influence item nonresponse in the mail survey. There are two exceptions to this conclusion. First, if respondents are shown several survey invitations that vary topic sensitivity and the disclosure risk associated with each one, the risk variable becomes statistically significant. Second, if

respondents are informed about the harm likely to result from disclosure, the combined risk-harm variable is statistically significant even if they are shown only one vignette. However, if, as in the mail survey, we separate the effect of risk and harm, we find that harm is significant whereas risk of disclosure is not. Neither of these conditions—that is, the presentation of varying risks under different scenarios, or the explicit spelling out of harms that might result from disclosure—is likely to become part of ordinary survey practice.

There is one other possible exception to our conclusion that we can only infer; we have no direct evidence for it. The strongest finding emerging from our various studies is the effect of survey topic. Both our sensitive topics—sexual and financial behavior—have significantly lower participation rates than the nonsensitive topics of leisure activities and work, regardless of whether they are described in hypothetical vignettes or actual cover letters for survey questionnaires. Invasions of privacy loom large in the minds of a majority of respondents who say they would refuse to participate in the hypothetical survey described, and among those who returned the nonresponse postcard to the mail survey, almost five times as many of those who had received questionnaires on sensitive topics checked “topic too sensitive” as among those who had received questionnaires on non-sensitive topics.

Thus, it is likely that sensitive topics themselves suggest to respondents the possibility (risk) of more harmful consequences, thus rendering explicit risk and harm information redundant. This may be especially true of respondents who have engaged in normatively deviant or illegal behavior. The fact that we find no significant interaction between topic sensitivity and disclosure risk in the mail survey may simply mean that too few respondents fall into this category to demonstrate statistical significance.

The conclusions above are, of course, limited by the characteristics of the mail survey. Perhaps the most important of these is the low return rate—28.3% of those exposed to the experimental conditions. Had the return rate been significantly higher, it is possible that an effect of disclosure risk might have emerged. In any event, we cannot rule out the possibility that a survey on a different topic, with a different sponsor and different financial incentives, would have produced different results with respect to the effects of disclosure risk.

The conclusions may also be limited by the definitions of harm we included in the web and mail surveys. However, we have now experimented with disclosure of harm defined in various ways in a telephone survey, several web surveys, and a mail survey, and we feel fairly confident that the negative effect of such a disclosure on participation is robust across various modes and definitions.

Despite these limitations, we conclude on the basis of all our experiments so far that specific mention of the size of the disclosure risk associated with a survey is unlikely to be used by respondents in deciding whether or not to participate in that survey. While it is important to estimate such risks and locate their sources in order to guard against them, these estimates should not be included in informed consent statements presented to potential respondents. Instead, researchers must find other means of assuring respondents that all possible measures to safeguard the confidentiality of their answers will be taken, without providing unrealistic estimates of such protection. The best ways of doing this remain to be demonstrated.

## 5. References

- Conrad, F.G., Park, H., Singer, E., Couper, M.P., Hubbard, F., and Groves, R.M. (2006). Impact of Disclosure Risk on Survey Participation Decisions. Paper presented at the annual meeting of the American Association for Public Opinion Research, Montreal, Canada, May.
- Couper, M.P., Singer, E., Conrad, F.G., and Groves, R.M. (2008). Risk of Disclosure, Perceptions of Risk, and Concerns about Privacy and Confidentiality as Factors in Survey Participation. *Journal of Official Statistics*, 24, 255–275.
- Couper, M.P. and Singer, E. (2009). The Role of Numeracy in Informed Consent for Surveys. *Journal of Empirical Research on Human Research Ethics*, 4, 17–26.
- Kish, L. (1987). *Statistical Design for Research*. New York: John Wiley.
- Malin, B. and Sweeney, L. (2000). Determining the Identifiability of DNA Database Entries. *Proceedings of the Journal of the American Medical Informatics Association*. Washington, DC: Hanley and Belfus, Inc, November, 537–541.
- Paass, G. (1988). Disclosure Risk and Disclosure Avoidance for Microdata. *Journal of Business and Economic Statistics*, 6, 487–500.
- Raghunathan, T.E., Lepkowski, J.M., Van Hoewyk, J., and Solenberger, P. (2001). A Multivariate Technique for Multiply Imputing Missing Values Using a Sequence of Regression Models. *Survey Methodology*, 27, 85–95.
- Raghunathan, T.E., Lepkowski, J.M., Van Hoewyk, J., and Solenberger, P. (2005). IVEware, A Software Package for the Analysis of Complex Survey Data with or Without Multiple Imputations. University of Michigan, Institute for Social Research, ([www.isr.umich.edu/src/smp/ive](http://www.isr.umich.edu/src/smp/ive))
- Sherif, M. (1958). Group Influences upon the Formation of Norms and Attitudes. In *Readings in Social Psychology*, E.E. Maccoby, T.M. Newcomb, and E.L. Hartley (eds). (3rd Edition). New York: Henry Holt and Co., 219–232.
- Singer, E. (2003). Exploring the Meaning of Consent: Participation in Research and Beliefs About Risks and Benefits. *Journal of Official Statistics*, 19, 273–286.
- Winkler, W.E. (1997). Views on the Production and Use of Confidential Microdata. U.S. Census Bureau Research Report No. RR97/01.

Received July 2008

Revised February 2010