



^b
**UNIVERSITÄT
BERN**

Faculty of Business, Economics and
Social Sciences

Department of Social Sciences

University of Bern Social Sciences Working Paper No. 18

More Is Not Always Better: An Experimental Individual-Level Validation of the Randomized Response Technique and the Crosswise Model

Marc Höglinger and Ben Jann

February 15, 2016

<http://ideas.repec.org/p/bss/wpaper/18.html>
<http://econpapers.repec.org/paper/bsswpaper/18.htm>

**More Is Not Always Better: An Experimental Individual-Level Validation
of the Randomized Response Technique and the Crosswise Model**

Marc Hoeglinger

ETH Zurich
Clausiusstrasse 50
CH-8092 Zurich

marc.hoeglinger@soz.gess.ethz.ch

Ben Jann

University of Bern
Institute of Sociology
Fabrikstrasse 8, CH-3012 Bern

ben.jann@soz.unibe.ch

February 2016

Acknowledgments: The authors thank Andreas Diekmann for his support and advice, Philip Tschiemer for his help with programming the survey, Debra Hevenstone for language editing, and Stefan Wehrli and the ETH Decision Science Laboratory (DeSciL) for posting the survey on Amazon Mechanical Turk.

Funding: This study was funded by the German Research Foundation (DFG) as part of the project “Asking Sensitive Questions: Possibilities and Limits of Randomized Response and Other Techniques in Different Survey Modes” (DI 292/5 to Andreas Diekmann), the Chair of Sociology of the ETH Zurich, and the Institute of Sociology of the University of Bern.

Abstract

Social desirability and the fear of sanctions can deter survey respondents from responding truthfully to sensitive questions. Self-reports on norm breaking behavior such as shoplifting, non-voting, or tax evasion may therefore be subject to considerable misreporting. To mitigate such misreporting, various indirect techniques for asking sensitive questions, such as the randomized response technique (RRT), have been proposed in the literature. In our study, we evaluate the viability of several variants of the RRT, including the recently proposed crosswise-model RRT, by comparing respondents' self-reports on cheating in dice games to actual cheating behavior, thereby distinguishing between false negatives (underreporting) and false positives (overreporting). The study has been implemented as an online survey on Amazon Mechanical Turk ($N = 6,505$). Our results indicate that the forced-response RRT and the unrelated-question RRT, as implemented in our survey, fail to reduce the level of misreporting compared to conventional direct questioning. For the crosswise-model RRT, we do observe a reduction of false negatives (that is, an increase in the proportion of cheaters who admit having cheated). At the same time, however, there is an increase in false positives (that is, an increase in non-cheaters who falsely admit having cheated). Overall, our findings suggest that none of the implemented sensitive questions techniques substantially outperforms direct questioning. Furthermore, our study demonstrates the importance of distinguishing false negatives and false positives when evaluating the validity of sensitive question techniques.

Keywords: Sensitive Questions, Online Survey, Amazon Mechanical Turk, Randomized Response Technique, Crosswise Model, Dice Game, Validation

1 Introduction

Surveying sensitive topics such as deviant behavior, stigmatizing traits, or controversial attitudes poses serious challenges to survey research. First, respondents' data need to be carefully protected, particularly for sensitive themes like illegal behavior or politically repressed opinions. Second, even with good data protection, respondents might be tempted to misreport on sensitive questions or refuse to answer, for example, due to embarrassment or due to fear of negative sanctions (Tourangeau and Yan 2007). To avoid biased or incomplete measurement, survey researchers therefore have to find questioning procedures that maximize respondents' willingness to provide truthful answers.

Various approaches to address this issue have been pursued in previous research, but the results on the success and the failure of the different questioning strategies appear inconsistent and highly dependent on implementation details, the research question, or the studied population (Krumpal and Näher 2012). Most promising results can be found with respect to survey mode and, in particular, to whether an interviewer is present or not. For example, Kreuter et al. (2008) compared CATI (Computer Assisted Telephone Interviewing), IVR (Interactive Voice Response), and online mode in a study on poor (and potentially embarrassing) academic performance among university alumni, where the respondents' answers could be validated against the university's grade records. The level of misreporting (false denial of poor performance) was highest in CATI mode, where an interviewer was present. However, also in the more anonymous IVR and online modes, misreporting remained high.

The randomized response technique

Other approaches try to mitigate misreporting and non-response by employing so-called indirect questioning techniques, one of which is the randomized response technique (RRT; originated by Warner 1965). The basic idea of the RRT is to protect respondents through random misclassification so that a given answer does not reveal the true answer to the sensitive question. Ideally the anonymity induced by the misclassification makes respondents more comfortable providing truthful answers. For example, in the forced-response variant of the RRT (Boruch 1971) a randomizing device such as a coin flip determines whether a respondent is instructed to provide a truthful answer to a sensitive yes/no question or simply respond with “yes” (or “no”), irrespective of the true answer. Therefore, as long as only the respondent knows the outcome of the randomizing device, a given answer does not reveal the true answer to the sensitive question; the given answer could also just be a surrogate response due to the randomizing device.

Despite the theoretical appeal of the RRT, it remains questionable whether respondents understand the procedure, trust that their anonymity is protected, and are more inclined to provide a truthful answer (when instructed to do so). Furthermore, due to lack of understanding, respondents might fail to comply with the RRT instructions even if they are asked to provide an answer that is unrelated to the sensitive question (Edgell et al. 1982; Edgell et al. 1992; Böckenholt et al. 2009). A meta analysis by Lensvelt-Mulders et al. (2005), mostly covering face-to-face and paper-and-pencil RRT studies published between 1965 and 2000, concludes that, on average, the RRT yields more valid results than direct questioning, but the variability in results is high. Furthermore, findings from a number of newer studies on the application of the RRT in online mode are not very promising (Coutts

and Jann 2011; Höglinger et al. 2014; Holbrook and Krosnick 2010; Ostapczuk und Musch 2011; Peeters 2005).

The crosswise-model RRT

Recently, a variant of the RRT, the “crosswise model,” proposed by Yu et al. (2008), has received growing attention. Several studies report that the crosswise-model RRT consistently produces higher prevalence estimates of sensitive behaviors than direct questioning (Gingerich et al. 2015; Hoffmann et al. 2015; Hoffmann and Musch 2015; Höglinger et al. 2014; Jann et al. 2012; Korndörfer et al. 2014; Kundt 2014; Kundt et al. 2014; Shamsipour et al. 2014). The crosswise-model RRT works by presenting two yes/no questions to the respondent, a sensitive question and an unrelated non-sensitive question, and then asking whether the answers to both questions are the same (both “yes” or both “no”) or whether the two answers are different (one “yes,” one “no”). The advantages of the crosswise-model RRT over alternative RRT variants, it is argued, are that the instructions are easy to understand, the response options are obviously ambiguous with respect to the sensitive question (i.e. there is no clear self-protective answering strategy), and no respondents are forced to give “false” answers.

Validation of sensitive question techniques

As mentioned above, results from studies evaluating indirect questioning techniques are often inconclusive. One reason for the variability in the findings is that the studies employ different validation strategies. By far the most frequent approach is to use the results from direct questioning as a baseline, to which the results from one or several indirect questioning

techniques are compared. We use the term *comparative validation study* to refer to studies employing such an approach. The argument is that if the question is sensitive, respondents will tend to underreport when asked to answer the question directly. An indirect questioning technique that successfully reduces underreporting should therefore yield higher estimates than direct questioning (likewise, if the problem is over-reporting, such as in questions on voter turnout, a successful indirect technique should yield lower estimates than direct questioning). Hence, comparative validation studies rely on the so-called more-is-better (less-is-better) assumption (Lensvelt-Mulders et al. 2005); an indirect questioning technique is considered more valid if it produces higher (lower) prevalence estimates than direct questioning. More generally, if comparing multiple indirect techniques, the technique producing the highest (lowest) estimate is judged to be the most valid.

The more-is-better assumption is often legitimate. In many cases it is reasonable to assume that respondents avoid socially undesirable answers and thus underreport on sensitive questions. However, sometimes, social desirability might differ between subpopulations, a well-known example being the number of sexual partners as reported by men and women (Smith 1992; Tourangeau and Smith 1996). Therefore, the more-is-better assumption can sometimes be challenged on the ground that social desirability bias points in different directions depending on the subpopulation. Furthermore, even if the more-is-better assumption is justified, a higher estimate from an indirect questioning technique does not necessarily imply that the technique produces more valid measurements than direct questioning. If it is true that direct questioning yields underestimation, then higher estimates by an indirect technique is a necessary condition, but not a sufficient condition. The more-is-better assumption assumes that the increase in estimates is due to more truthful answers. However, given the complexity of the instructions of most RRT implementations it may also

be simply due to the respondents' inability to correctly apply the procedure. That is, the increase in estimates might be due to non-compliance with the RRT instructions (e.g., due to problems with the randomizing device, misunderstanding of instructions, or unwillingness to follow the instructions) rather than more truthful answering. Overall, we conclude that comparative validation studies can only provide weak support for the validity of sensitive questioning techniques (for similar arguments see: Lensvelt-Mulders et al. 2005; Höglinger et al. 2014; Moshagen et al. 2014; Wolter and Preisendörfer 2013).

At least some of the shortcomings of comparative validation studies can be overcome by what we call *aggregate-level validation studies*. In such studies, the true population prevalence of the sensitive trait or behavior is known from an external and reliable source or can be determined based on theoretical reasoning. For example, in studies of voter turnout, true aggregate turnout is known from administrative records (for recent examples see Rosenfeld et al. 2015 and Moshagen et al. 2012). If the true value is known, then overestimation and underestimation by different questioning techniques can be observed directly without having to resort to direct questioning as a baseline, which is a clear improvement over comparative validation studies. Yet, also such aggregate-level validation studies might be inconclusive. First, true values might differ from the assumed value, perhaps because the study focuses on a special subpopulation or because there is sample selection bias (e.g., due to nonresponse). Second, and more importantly, a close match between the prevalence estimate from a particular sensitive questioning technique and the true value does not necessarily imply that the technique produces valid measurements at the individual level. As argued above, different mechanisms might affect the prevalence estimate, not all of which are consistent with more truthful answering. In other words, apart from possible sample selection bias, the aggregate-

level validation approach rests on the assumption that socially desirable responding is the only misreporting mechanism.

A useful distinction in this context is between false negatives (or true positives) and false positives (or true negatives). The goal of sensitive question techniques is to reduce the number of false negatives, that is, the number of respondents who deny the sensitive question even though it does apply. However, a sensitive question technique might also increase the number of false positives, that is, the number of respondents who agree with the sensitive question even though it does not apply. Comparing overall prevalence estimates from the technique with either direct questioning or a known “true” prevalence, does not allow one to distinguish between a reduction in false negatives and an increase in false positives, both of which will increase the estimated total prevalence. To be able to disentangle the two effects, validation data at the individual level is required. Hence, we argue that *individual-level validation studies* are necessary to be able to evaluate the degree to which a technique does, in fact, produce valid measurements.

Despite their clear advantage over the comparative approach, individual-level validation studies are very rare. Reviewing RRT studies from over 35 years, Lensvelt-Mulders et al. (2005) counted just six published individual-level validation studies dealing with sensitive topics such as convictions, arrests, welfare fraud, or failing university courses. We are aware of five additional studies published since (Hoffmann et al. 2015; John et al. 2013; Kirchner 2015; Moshagen et al. 2014; Wolter and Preisendörfer 2013). The available validation studies provide valuable insights, but they do not explicitly focus on disentangling false negatives and false positives. Moreover, some of the studies use a sample that only includes respondents who possess the sensitive trait or engaged in the sensitive activity, so that, by

design, only false negatives can be studied. In sum, we believe that additional individual-level validation studies are necessary to disentangle the different response mechanisms and to examine the possibility of false positives in these types of survey techniques. Such studies are the only way to conclusively assess the performance of different sensitive question techniques.

Our study

The goal of our study is to evaluate the validity of different variants of the RRT using a validation design that does not rely on the more-is-better assumption and that allows separate analysis of false negative and false positives. To achieve this we conducted an online survey on Amazon Mechanical Turk, in which the respondents had the opportunity to play one of two dice games. Respondents were given monetary incentives to cheat in these games. After playing the games, respondents were asked about whether they cheated, using direct questioning, forced-response RRT, unrelated-question RRT, or the crosswise-model RRT. For the first game the proportion of cheaters can be estimated based on the laws of chance, for the second game cheating is observable. Comparing the cheating behavior at the aggregate and individual levels with the results from the cheating question reveals the degree to which the questioning techniques are successful in eliciting truthful answers.

2 Data and Methods

Study participants were recruited via the online platform Amazon Mechanical Turk (AMT). AMT is an online crowdsourcing marketplace where “requesters” can post tasks (called “Human Intelligence Tasks” or HITs) that can then be completed by “workers” in exchange

for money. HITs are announced with a short description of the task and the corresponding payment. AMT is suitable for any task that can be easily outsourced online to an anonymous workforce and is increasingly used to recruit participants for scientific surveys and experiments (Horton et al. 2011; Mason and Suri 2012; Ipeirotis 2010). On November 5, 2013, we posted a HIT that asked for filling out a scientific survey on “Mood and Personality” for a base payment of \$1 and the prospect of winning an additional \$2 bonus payment. The HIT was closed on December 5, 2013, when our quotas per experimental condition were fulfilled. Workers who accepted our HIT received an access link to the survey. After having completed the survey, they received payment. Participation was restricted to US residents because one of the sensitive questions was on voting in the US presidential elections. To identify untrustworthy participants, we employed a screening question from Berinsky et al. (2012), which was passed by 97% of the respondents. The median time required to complete the survey was 6.7 minutes. Details on study and screenshots of the questionnaire are available in the survey documentation provided in the online appendix.

A total of 6,505 participants were recruited, of which 6,473 completed the survey at least up to the part containing the sensitive questions. Only the latter are included in our analysis. Furthermore, we exclude 205 participants who did not pass the screening question, 115 participants who did not roll the die in the dice game (or for whom the result of the roll was not recorded due to technical problems), and 1 participant who won in the roll-a-six game but did not claim his legitimate bonus payment.¹ The final sample size for our analysis is $N = 6'152$. As displayed in table 1, the sample has an even gender distribution and the majority of

¹ There were 516 winners in the roll-a-six game. The fact that only one of them did not claim the bonus payment indicates that the proportion respondents who falsely deny having won is negligible. To simplify the analysis we exclude this observation and assume the proportion to be zero (also in the prediction game, where winners cannot be identified at the individual level).

respondents are under 35 (mean age 32). Respondents are relatively well educated, with 88 percent having attended at least some college. About two thirds are employed or self-employed. A large majority of respondents completed the survey at home and most respondents had extensive experience with “scientific studies such as surveys or experiments on MTurk” (wording from the questionnaire; the median number of previous MTurk studies is 50).

[table 1 about here]

The dice games

Participants were randomly assigned to one of two dice games in which they could win a \$2 bonus payment: the *prediction game* or the *roll-a-six game*. The games were inspired by Greene and Paxton (2009) and Fischbacher and Föllmi-Heusi (2013) (also see Fischbacher and Heusi 2008 as well as Suri et al. 2011). In both games, participants used a digital online die embedded in the questionnaire that could be “rolled” by clicking on a button. Roll outcomes were randomized and followed a uniform distribution. The die could be rolled several times, but as explained to the respondents, only the first roll counted.

In the *prediction game* participants had to correctly predict the outcome of a die roll to win the \$2 bonus payment. On a first screen, the rules of the game and the conditions under which a participant would win the bonus payment were explained. On the second screen, participants were asked to make their prediction (in private) and memorize it. On the third screen they were instructed to roll the die, inspect the result, and then indicate whether their prediction was correct or not. Because the prediction was made in private, cheating could not

be detected. Since the probability of winning was one sixth, however, the proportion of cheating respondents can be estimated at the aggregate level (assuming that all respondents whose prediction was correct do claim the bonus payment).

In the *roll-a-six game* participants had to roll a six in order to win the \$2 bonus payment. Respondents were again presented a first screen on which the game was explained. On the second screen they were instructed to roll the die and then indicate whether the result was a six or not. As in the prediction game, cheating was easily possible as the bonus payment was determined solely on the basis of the respondent's answer and not on the actual outcome of the roll. Furthermore, estimation of the proportion of cheaters is again possible at the aggregate level as the theoretical probability of winning was one sixth. In contrast to the prediction game, however, also the identification of individual cheaters is possible since the outcomes of the die roll were recorded. Although respondents were not told that the outcomes would be tracked, it was clear that this was possible. Therefore, the proportion of cheaters can be expected to be lower in the roll-a-six game than in the prediction game. Likewise, when asked about whether they cheated, cheating respondents in the roll-a-six game may be expected to provide more truthful answers than cheating respondents in the prediction game.

The sensitive question techniques

In the second part of the questionnaire, respondents were asked four "sensitive" questions, the last of which being about whether they gave an honest answer in the dice game (see table 2).

[table 2 about here]

To evaluate different sensitive question techniques, respondents were randomly assigned to one of four conditions: direct questioning (DQ), the crosswise-model RRT (CM), the unrelated-question RRT (UQ), or the forced-response RRT (FR). Table 3 reports the number of observations per sensitive question technique and dice game variant. Respondents were unevenly distributed across conditions in order to counterbalance the different statistical efficiencies of the procedures.²

[table 3 about here]

Direct questioning (DQ) was included as a benchmark for the evaluation of the different sensitive question techniques. The sensitive questions were introduced by a screen announcing some sensitive questions, stating the importance of honest answers for the success of the study, providing privacy assurance, and telling the respondents that their answers to the sensitive questions would not affect their payment or the HIT approval (this introductory screen was identical for all conditions). After that, the four sensitive questions followed on four separate screens.

For the crosswise-model RRT (CM) we used an implementation as proposed in Jann et al. (2012) and Höglinger et al. (2014). Respondents were asked two questions: A sensitive question and an unrelated non-sensitive question. Respondents then had to indicate whether their answers to the two questions were the same (both “no” or both “yes”) or different (one “yes,” one “no”) without reporting the individual answers. The unrelated questions, which were randomly paired with the sensitive questions for each respondent, asked about the

² Item-nonresponse was negligible; below 1% for all sensitive questions in all experimental conditions. We therefore refrain from reporting results on item-nonresponse in the analyses below.

birthday (in January or February, between the 1st and the 6th of the month) of the respondent's mother or father. Between the introductory screen and the screen with the first sensitive question, an additional screen was displayed explaining the questioning technique and how it protects anonymity (similar screens were also displayed for the other sensitive question techniques).

For the unrelated-question RRT (UQ) we used an implementation as proposed by Diekmann (2012). Respondents were asked to think of an acquaintance and use the first digit of this person's house number as their personal random number. If their random digit was 1, 2, 3, 4, or 5, respondents then had to answer the subsequent sensitive questions; otherwise they had to answer the subsequent unrelated non-sensitive questions. Diekmann (2012) provides evidence that first digits of house numbers follow "Benford's Law". Accordingly, the probability of 1, 2, 3, 4, or 5 (i.e., of having to answer the sensitive questions) is 0.778.³ The unrelated questions were randomly paired with the sensitive questions for each respondent and asked about the birthday of the respondent's mother (in January–June, in an even-numbered month, in the first half of the month, on an even-numbered day, in an even-numbered year).

For the forced-response RRT (FR) we used an implementation as proposed by Höglinger et al. (2014). Respondents were presented twelve fields on the screen, numbered from one to twelve. They were told to privately choose a field and memorize their choice (without

³ To evaluate whether Benford's Law holds, we included a question on the first digit of an acquaintance's address for a subsample of respondents in a different experimental condition. The proportion of respondents reporting a 1, 2, 3, 4, or 5 was 0.784 (95% confidence interval: 0.763 to 0.804). Similar tests were included for all unrelated questions used in CM and UQ. Since deviations between the theoretical values (assuming an even distribution of birthdays) and the estimated proportions were only small, we focus on results based on the theoretical values in the analyses below.

clicking on the field). Then, they were told to click a “Show instructions” button to uncover the instructions hidden within the fields and follow the instruction that appeared in the field of their choice. Possible instructions were “Answer question”, “Directly tick yes”, or “Directly tick no”. The instructions were randomized across fields.

Data analysis

The RRT leads to data misclassification so that adjusted methods for data analysis are required. Let Y^* be the (unobserved) answer to the sensitive question ($Y^* = 1$ if the answer “yes”, $Y^* = 0$ else) and Y be the observed response ($Y = 1$ if the response is “yes” in case of DQ, UQ and FR or “the same” in case of CM; $Y = 0$ else).⁴ For direct questioning, $Y = Y^*$. The RRT procedures, however, introduce misclassification so that $Y \neq Y^*$. In general, in a misclassification setting, the relation between Y and Y^* can be described as

$$\Pr(Y = 1) = \Pr(Y = 1|Y^* = 1) \Pr(Y^* = 1) + \Pr(Y = 1|Y^* = 0) \Pr(Y^* = 0)$$

Solving for $\Pr(Y^* = 1)$ yields

$$\Pr(Y^* = 1) = \lambda(\Pr(Y = 1)) = \frac{\Pr(Y = 1) - p_{1|0}}{p_{1|1} - p_{1|0}}$$

⁴ Throughout this discussion we assume that “yes” is the sensitive answer, although some of the sensitive questions in our study were framed differently (for example, we asked respondents whether played honestly in the dice game, not whether they cheated). For the purpose of analysis, all data was appropriately recoded.

with $p_{1|1} = \Pr(Y = 1|Y^* = 1)$ and $p_{1|0} = \Pr(Y = 1|Y^* = 0)$. In the RRT, $p_{1|1}$ and $p_{1|0}$ are known by design. Hence, we can estimate $\Pr(Y^* = 1)$ by inserting a sample estimate for $\Pr(Y = 1)$ (i.e. the sample mean \bar{Y}) into the above formula. Furthermore, since $\Pr(Y^* = 1)$ is a linear transformation of $\Pr(Y = 1)$ and, in general, $V(ax + b) = a^2V(x)$ (see, e.g., Mood et al. 1974: 179), the sampling variance of estimator $\widehat{\Pr}(Y^* = 1)$ is given as

$$V(\widehat{\Pr}(Y^* = 1)) = \frac{1}{(p_{1|1} - p_{1|0})^2} V(\widehat{\Pr}(Y = 1))$$

where $V(\widehat{\Pr}(Y = 1))$ can be estimated from the data using standard techniques (e.g. as $\bar{Y}(1 - \bar{Y})/(n - 1)$ where n is the sample size). For direct questioning, there is no misclassification, so that $p_{1|1} = 1$ and $p_{1|0} = 0$ and hence

$$\lambda(\Pr(Y = 1)) = \Pr(Y = 1)$$

For the CM, let p_Z be the known probability that the answer to the non-sensitive question is “yes”. Then $p_{1|1} = p_Z$ and $p_{1|0} = 1 - p_Z$. Hence,

$$\lambda(\Pr(Y = 1)) = \frac{\Pr(Y = 1) + p_Z - 1}{2p_Z - 1}$$

For UQ, again let p_Z be the known probability that the answer to the non-sensitive question is “yes.” Furthermore, let p_U be the probability that the respondent is instructed to answer the non-sensitive question instead of the sensitive question. We then have $p_{1|1} = 1 - p_U(1 - p_Z)$ and $p_{1|0} = p_U p_Z$, so that

$$\lambda(\Pr(Y = 1)) = \frac{\Pr(Y = 1) - p_U p_Z}{1 - p_U}$$

Finally, for FR, let p_{yes} and p_{no} be the probabilities of an unconditional “yes” or “no” answer, respectively. Then $p_{1|1} = 1 - p_{\text{no}}$ and $p_{1|0} = p_{\text{yes}}$, so that

$$\lambda(\Pr(Y = 1)) = \frac{\Pr(Y = 1) - p_{\text{yes}}}{1 - p_{\text{yes}} - p_{\text{no}}}$$

The above formulas can be used to obtain prevalence estimates for the sensitive behaviors.

Employing the more-is-better assumption or comparing the estimates to the aggregate cheating rates in the dice games, we can then decide which of the techniques works best. The formulas, however, assume that respondents comply with the instructions so that, for example, no false positives occur (apart from false positives induced by design). If this assumption is violated, then the overall estimates can be misleading. To evaluate the degree to which the techniques produce valid results, we therefore perform separate analyses for those who cheated in the dice game and for those who did not cheat. What we are interested in is the *true positive rate* (TPR) that is, the proportion of cheaters who admit having cheated, and the *false positive rate* (FPR), that is, the proportion of non-cheaters who falsely “admit” having cheated. Furthermore, as overall measure of validity, we are interested in the *correct classification rate* (CCR).

For the roll-a-six game, these analyses are straightforward since cheating is observed at the individual level. Let $X^* = 1$ if the respondent rolled a six and $X^* = 0$ else. Furthermore, let $X = 1$ if the respondent *claimed* having rolled a six and $X = 0$ else. A respondent is identified

as a cheater (false winner) if $X = 1$ even though $X^* = 0$. Non-cheaters are given if $X = X^*$, that is, if $X = X^* = 1$ (true winner) or $X = X^* = 0$ (true loser). For sake of simplicity, assume that “reverse” cheating ($X = 0$ even though $X^* = 1$) is nonexistent, that is, assume that there are no respondents who did roll a six but then did not claim the bonus payment (false losers). The *true positive rate* is then given as

$$\text{TPR} = \Pr(Y^* = 1|X \neq X^*) = \lambda(\Pr(Y = 1|X \neq X^*))$$

and the *false positive rate* is give as

$$\text{FPR} = \Pr(Y^* = 1|X = X^*) = \lambda(\Pr(Y = 1|X = X^*))$$

Furthermore, the *correct classification rate* is

$$\text{CCR} = \text{TPR} \cdot \Pr(X \neq X^*) + (1 - \text{FPR}) \Pr(X = X^*)$$

Since X^* is observed in the roll-a-six game, all of the above quantities can be readily estimated from the data. In the prediction game, however, X^* is unobserved (in the prediction game, X^* denotes whether the respondent’s prediction was correct or not, X denotes whether the respondent *claimed* that the prediction was correct). Again, assume that all respondents whose predictions were correct do claim the bonus payment (no false losers), that is, that the combination $X^* = 1$ and $X = 0$ does not exist (as mentioned above, only one of 516 winners in the roll-a-six game did not claim the bonus payment; it appears highly plausible to assume that the proportion of false losers is negligible also in the prediction game). We know from

the design of the game that $\Pr(X^* = 1) = 1/6$, so that the proportion of cheaters, given that there are no false losers, is equal to

$$\Pr(X \neq X^*) = \Pr(X = 1) - \Pr(X^* = 1) = \Pr(X = 1) - 1/6$$

Furthermore, the false positive rate of true losers is given as

$$\Pr(Y^* = 1|X = X^* = 0) = \lambda(\Pr(Y = 1|X = 0))$$

The overall false positive rate or the true positive rate, however, cannot be identified without further assumptions. In general, the true positive rate can be written as

$$\Pr(Y^* = 1|X \neq X^*) = \frac{\Pr(Y^* = 1 \cap X \neq X^*)}{\Pr(X \neq X^*)} = \frac{\Pr(Y^* = 1 \cap X \neq X^*)}{\Pr(X = 1) - 1/6}$$

To identify $\Pr(Y^* = 1 \cap X \neq X^*)$ we need to make an assumption about the false positive rate of winners. The most reasonable assumption, in our opinion, is that the false positive rate of winners is equal to the false positive rate of true losers. Both types of respondents were honest in the prediction game and we do not see much reason why they should differ in their response behavior when asked about whether they were honest or not.⁵ That is, we assume $\Pr(Y^* = 1|X = X^* = 1) = \Pr(Y^* = 1|X = X^* = 0)$ and, hence, $\Pr(Y^* = 1|X = X^*) = \Pr(Y^* = 1|X = X^* = 0)$, so that the overall *false positive rate* can be written as

⁵ Note, however, that the composition of the two groups is somewhat different. Among the winners there are *potential* cheaters, that is, respondents who would have cheated should they not have won, as well as non-cheaters. The group of true losers only contains non-cheaters. Differential assumptions about the response behavior of potential cheaters and non-cheaters could be made, but would not fundamentally change our results.

$$\text{FPR} = \Pr(Y^* = 1|X = X^*) = \lambda(\Pr(Y = 1|X = 0))$$

For the derivation of the true positive rate note that

$$\Pr(Y^* = 1 \cap X \neq X^*) = \Pr(Y^* = 1 \cap X = 1) - \Pr(Y^* = 1 \cap X = X^* = 1)$$

(again given that there are no false losers). Since

$$\Pr(Y^* = 1 \cap X = 1) = \Pr(X = 1) \Pr(Y^* = 1|X = 1) = \Pr(X = 1) \lambda(\Pr(Y = 1|X = 1))$$

and, from results from above,

$$\begin{aligned} \Pr(Y^* = 1 \cap X = X^* = 1) &= \Pr(X = X^* = 1) \Pr(Y^* = 1|X = X^* = 1) \\ &= \Pr(X^* = 1) \Pr(Y^* = 1|X = X^*) = \frac{1}{6} \lambda(\Pr(Y = 1|X = 0)) \end{aligned}$$

the *true positive rate* is given as

$$\text{TPR} = \Pr(Y^* = 1|X \neq X^*) = \frac{\Pr(X = 1) \lambda(\Pr(Y = 1|X = 1)) - \lambda(\Pr(Y = 1|X = 0))/6}{\Pr(X = 1) - 1/6}$$

Furthermore, the *correct classification rate* is

$$\text{CCR} = \text{TPR} \cdot (\Pr(X = 1) - 1/6) + (1 - \text{FPR})(\Pr(X = 0) + 1/6)$$

3 Results

Comparative validation

We first report results as in a standard comparative validation study, using the more-is-better assumption. Figure 1 displays the point estimates for the sensitive behaviors from the different sensitive question techniques, as well as the differences in the estimates between direct questioning (DQ) and the indirect techniques (also see table A1 in the appendix). For shoplifting, estimates from all three indirect techniques are significantly higher than the estimate from direct questioning. The highest estimate was obtained by the unrelated-question RRT (UQ). Also for tax evasion, all three techniques significantly outperformed direct questioning, with the crosswise-model RRT (CM) producing the highest estimate. CM also produced the highest estimates for the remaining three items, although the difference to direct questioning is not significant for the non-voting item. The unrelated-question RRT (UQ) and the forced-response RRT (FR) did not produce significantly higher estimates than direct questioning for these three items. Moreover, for cheating in the roll-a-six game, the estimate from FR is significantly lower than the estimate from DQ. From these results we would conclude that the CM clearly performed best of all techniques; it produced the highest estimates for four of the five items and produced significantly higher estimates than direct questioning for the four of the five items. The difference between CM and the other techniques is particularly pronounced for the two cheating items. While cheating rates were 5% or less according to the other techniques, they were about 15% according to CM. The results for UQ and FR are mixed. They outperformed direct questioning for the first two

items, but not for the remaining three. For the last item, FR even produced a slightly negative estimate, indicating substantial non-compliance with the RRT instructions.⁶

[figure 1 about here]

Aggregate-level validation

As illustrated above, were we to conduct a comparative validation study based on the more-is-better assumption, we would find that the crosswise-model RRT is the most valid technique. However, the more-is-better assumption is a strong assumption that might be violated. In the second step, we therefore compare the prevalence estimates from the various techniques to the true prevalence of the sensitive behaviors at the aggregate level. We can conduct such an analysis for the two items on cheating in the dice games. Figure 2 displays the true rates as well as the various estimates including 95% confidence intervals (left panel).⁷ In the right panel of the figure, the differences between the true rates and the estimates are shown. For the prediction game, all questioning techniques performed poorly. DQ, UQ and FR all produced estimates below 5% although the true cheating rate was around 25%. The CM comes closest to the true cheating rate with an estimate of a bit more than 15%, but still underestimates the true rate by about 11 percentage points. For the roll-a-six game, we see that DQ and UQ both produced accurate estimates of a cheating rate of about 5%. As expected, cheating was

⁶ A negative estimate is possible if a substantial proportion of respondents deviate from the instructions determined by the randomizing device. This seems to be a common problem with the forced-response RRT (see, e.g., Coutts and Jann 2011).

⁷ Confidence intervals are also reported for the true cheating rates even though in the roll-a-six game the sample cheating rate can be determined exactly. The confidence intervals reflect the variability in the cheating rates one could expect were one to repeat the experiment.

substantially less prevalent in the roll-a-six game than in the prediction game, due to the design of the game (the roll-a-six game provided less incentive for cheating than the prediction game because it was obvious that cheating could potentially be detected). FR significantly underestimated the cheating rate. For CM, on the other hand, an overestimation by about 8 percentage points occurred. Hence, while for the prediction game the more-is-better assumption seems to be valid in the sense that the highest estimate comes closest to the true value, the assumption fails for the roll-a-six game. Respondents did not substantially underreport their cheating behavior in the roll-a-six game when asked directly, probably because it was obvious that such misreporting could be detected. One could argue that cheating in the roll-a-six game is therefore not a good test case for evaluating sensitive question techniques; there is no bias that could be improved on by the techniques. On the other hand, we would expect that a valid sensitive question technique produces unbiased results also if the question is, in fact, not sensitive. A positive bias such as observed for the CM should not occur.

[figure 2 about here]

Individual-level validation

Overall, the results from the aggregate-validation are ambiguous. For the first item, cheating in the prediction game, the crosswise-model RRT (CM) is the clear winner. If we had exclusively looked at the prediction game, we would have again concluded that CM is the most valid technique. However, cheating in the roll-a-six game indicates that there might be a problem with the CM. In the third step of our analysis we therefore evaluate the accuracy of the measurements obtained by the different questioning approaches at the individual level.

Figure 3 displays the true and false positive rates of the different techniques for the prediction game and the roll-a-six game. Direct questioning had a true positive rate (TPR) of only 10% in the prediction game, that is, only 10% of respondents who cheated in the prediction game admitted having done so when asked directly. FR did not manage to improve the TPR and UQ slightly increased the TPR to 15%. The CM, on the other hand, was considerably more successful in eliciting truthful answers from cheaters in the prediction game, with a true positive rate of almost 30% (although still being far from 100%). Yet, the CM also had a substantial false positive rate (FPR) of about 10%. That is to say, about 10% of respondents who did not cheat in the prediction game accidentally admitted having cheated when using the CM. Due to the (relatively) high TPR and the positive FPR the estimate of the cheating rate from the CM came closest to the true cheating rate at the aggregate level (as seen above). However, the correct classification rate (CCR) of the CM was, in fact, worst of all techniques (since about 75% of the respondents did not cheat, the positive FPR has a strong influence on the CCR). The UQ and FR did not have the problem of false positives, but did also not really improve on the TPR compared to DQ, so that these techniques did not reach a better CCR than DQ as well. Overall, for the prediction game, we can therefore conclude that the unrelated-question RRT (UQ) and the forced-response RRT (FR) did not manage to produce more accurate measurements than direct questioning, and that the crosswise-model RRT (CM), although seemingly more valid than direct questioning at aggregate level, fared worst in terms of correct classification at the individual level due to the occurrence of false positives.

[figure 3 about here]

For the roll-a-six game (right panel in figure 3) we obtain a similar picture. Also here the CM was affected by a substantial amount of false positives (to a similar degree as in the prediction game) and, again, although not severely affected by false positives, the UQ and FR did not perform better than direct questioning. For true positives, the ranking of the techniques changed in that direct questioning now performed best, with a true positive rate of about 70%. That the true positive rates for the indirect techniques were lower in this case than for direct questioning might be due to the fact that the RRT, although meant to provide an opportunity to be honest without the risk of disclosure, also provides respondents the possibility to be dishonest without the risk of disclosure. Because it was obvious in the roll-a-six game that a dishonest answer about whether a respondent cheated or not could potentially be identified, some of the respondents who would have felt compelled to answer truthfully in direct questioning might have misused the RRT as a protection mechanism to answer untruthfully without risk of detection.⁸ To summarize the results for the roll-a-six game: none of the indirect techniques managed to improve the true positive rate compared to direct questioning and the CM was affected by a substantial amount of false positives, so that similar to the prediction game, the correct classification was best for direct questioning and worst for the CM.

Our conclusion from the individual-level validation is that direct questioning, in fact, produced the most accurate measurements for both sensitive items. That is, from these results we have to conclude that direct questioning is the most valid technique. None of the tested indirect questioning techniques yielded an improvement over direct questioning. Keeping in

⁸ The possibility of such a paradoxical effect of indirect questioning techniques is also mentioned by Wolter and Preisendörfer (2013). Lelkes et al. (2012) found similar adverse effects of complete anonymity on truthful reporting.

mind that indirect techniques sacrifice statistical efficiency (and hence require larger sample sizes than direct questioning) we cannot recommend their general application (unless guaranteeing full privacy protection to respondents by misclassifying their answers is an important goal of a study). We also show that the CM is particularly problematic as it is affected by false positives. For example, the occurrence of false positives is the reason for why the CM overestimated the cheating rate in the roll-a-six game; it is also the reason why, at the aggregate level, the CM came seemingly closest to the true cheating rate in the prediction game. That the false positive rates of the CM were similar for both games indicates that there was a certain fraction of respondents in our sample who were unable or unwilling to apply the CM procedure correctly. How large this fraction is might depend on the specific population under study. It is clear, however, that the presence of such noncompliance has strong effects on the estimates obtained by the CM. We suspect that the false positives are the reason for why the CM performed so well in many previous studies that used a comparative design without the possibility for individual-level validation. False positives inflate the CM estimates and, from a more-is-better perspective, make it look like the CM provides more valid estimates than other techniques.

4 Conclusions

In order to evaluate the validity of survey respondents' self-reports based on various sensitive question techniques we carried out an online experiment in which respondents' self-reported rates of cheating were compared to true cheating rates. Participants played one of two incentivized dice games in which they could cheat, that is, in which they could illegitimately claim a bonus payment. After the game, participants were asked whether they cheated using

either direct questioning or one of several RRT implementations. The resulting self-reports were then validated against the actual rate of cheating in the dice game. Unlike most other evaluation studies of indirect questioning techniques, our study relies on a true validation criterion and detects misreporting at the individual level.

Results reveal that all tested questioning techniques suffer sizeable misclassification in the direction of the socially desirable answer. Among the different techniques only between 9% and 28% of all cheaters could be correctly classified as cheaters in the first variant of the dice game (prediction game). In the second variant of the dice game (roll-a-six game) between 41% and 71% of cheaters could be correctly classified. The large difference in the true positive rate between the two games suggests that the sensitivity of an item and – possibly, whether answers are potentially verifiable or not – has an important effect on respondents' decision whether to misreport or not. Although, at least for the prediction game, some of the evaluated indirect questioning techniques yielded higher true positive rates than direct questioning, none of the techniques produced overall more valid measurements than direct questioning. The reason is that the indirect techniques tend to produce poor results for respondents who do not possess the sensitive trait (i.e. who did not cheat). In particular, a substantial false positive rate was observed for the crosswise-model RRT (CM), that is, for the subsample of non-cheaters, the CM erroneously yielded cheating rates of about 11% or 12%. Furthermore, the forced-response RRT (FR) yielded negative cheating rates in the subsample of non-cheaters, which indicates that some of the respondents did not comply to the RRT instructions and answered “no” even though the procedure instructed them to answer “yes.” The unrelated-question RRT (UQ) had the least problems with respect to misclassification in the subsample of non-cheaters, but it did also not substantially reduce the amount of misclassification in the subsample of cheaters.

An important insight of our study is that the findings would have been quite different had there not been the possibility for individual-level validation. False positives in the CM inflated the prevalence estimates so that the CM consistently yielded higher prevalence of sensitive behaviors than direct questioning. Hence, employing the more-is-better assumption, the CM seemed superior. As illustrated by the first sensitive item in our study for which validation was possible (cheating in the prediction game), comparing prevalence estimates from indirect questioning techniques to the true prevalence rate at the aggregate level, although certainly an improvement over the more-is-better assumption, can still be misleading. The CM provided a prevalence estimate that came closest to the true prevalence. Hence, one could again conclude that the CM has superior validity. The analysis at the individual level, however, revealed that this is a false conclusion. The CM came close to the true prevalence primarily because it misclassified some of the non-cheating respondents as cheaters. That is, our study not only shows that the CM might not be as promising as suggested by previous studies (Gingerich et al. 2015; Hoffmann et al. 2015; Hoffmann and Musch 2015; Höglinger et al. 2014; Jann et al. 2012; Korndörfer et al. 2014; Kundt 2014; Kundt et al. 2014; Shamsipour et al. 2014), it also points to a general weakness in past research on sensitive question techniques. Because complicated misreporting patterns are possible, we must be very cautious when interpreting results from comparative evaluation studies employing the more-is-better assumption, from validation studies that rely on aggregated prevalence validation, or from one-sided validation studies in which the sensitive trait or behavior applies to all or none of the respondents. We argue that an integral evaluation of the performance of a sensitive questioning technique is only possible if answers can be validated at the individual level so that false negatives and false positives can be disentangled.

Of course, our study also has limitations. For example, we cannot answer why a substantial share of non-cheaters misreported in the CM. It is noteworthy that such misreporting did not occur with direct questioning. As such, we would speculate the cause might have to do with confusion rather than carelessness. It would be worthwhile to conduct further research on the CM to identify the design feature that causes this type of misreporting and to evaluate possible modifications to address the problem. Furthermore, our study uses two very specific items, cheating in the prediction game and cheating in the roll-a-six game, to evaluate the sensitive question techniques and, in addition, has been conducted in a special setting and in a special population (a survey on Amazon Mechanical Turk). Whether our results can be generalized to other sensitive questions, and to other populations and settings remains questionable. Further research should therefore investigate whether our results can be replicated in other contexts. Finally, we only evaluated three specific variants of the randomized response technique. Although the results of our study are discouraging for all three variants, there might be alternative designs or implementations that are more successful. Future research should focus on evaluating such alternatives. Using a research design that allows individual-level validation of respondents' answers, however, would be crucial for such research to be meaningful.

References

- Berinsky, Adam J., Michele Margolis, and Michael W. Sances. 2012. "Separating the Shirkers from the Workers? Making Sure Respondents Pay Attention on Internet Surveys." Paper presented at the NYU CESS 5th Annual Experimental Political Science Conference (available from http://cess.nyu.edu/policon2012/wp-content/uploads/2012/02/berinsky_margolis_sances_screener_nyu.pdf).
- Böckenholt, Ulf, Sema Barlas, and Peter G. M. van Der Heijden. 2009. "Do Randomized-Response Designs Eliminate Response Biases? An Empirical Study of Non-Compliance Behavior." *Journal of Applied Econometrics* 24:377-92.
- Boruch, Robert F. 1971. "Assuring Confidentiality of Responses in Social Research: A Note on Strategies." *The American Sociologist* 6:308-11.
- Coutts, Elisabeth, and Ben Jann. 2011. "Sensitive Questions in Online Surveys: Experimental Results for the Randomized Response Technique (RRT) and the Unmatched Count Technique (UCT)." *Sociological Methods & Research* 40:169-93.
- Diekmann, Andreas. 2012. "Making Use of 'Benford's Law' for the Randomized Response Technique." *Sociological Methods & Research* 41:325-34.
- Edgell, Stephen E., Karen L. Duchan, and Samuel Himmelfarb. 1992. "An Empirical Test of the Unrelated Question Randomized Response Technique." *Bulletin of the Psychonomic Society* 30:153-56.
- Edgell, Stephen E., Samuel Himmelfarb, and Karen L. Duchan. 1982. "Validity of Forced Responses in a Randomized-Response Model." *Sociological Methods & Research* 11:89-100.
- Fischbacher, Urs, and Franziska Föllmi-Heusi. 2013. "Lies in Disguise—An Experimental Study on Cheating." *Journal of the European Economic Association* 11:525-47.

- Fischbacher, Urs, and Franziska Heusi. 2008. "Lies in Disguise. An Experimental Study on Cheating." Research Paper Series No. 40, Thurgau Institute of Economics at the University of Konstanz (available from http://www.twi-kreuzlingen.ch/uploads/tx_cal/media/TWI-RPS-040-Fischbacher-Heusi-2008-11.pdf).
- Gingerich, Daniel, Virginia Oliveros, Ana Corbacho, and Mauricio Ruiz-Vega. 2015. "Corruption as a Self-Fulfilling Prophecy: Evidence from a Survey Experiment in Costa Rica." IDB Working Paper No. IDB-WP-546 (available from <https://publications.iadb.org/handle/11319/6826>).
- Greene, Joshua D., and Joseph M. Paxton. 2009. "Patterns of Neural Activity Associated with Honest and Dishonest Moral Decisions." *Proceedings of the National Academy of Sciences* 106:12506-11.
- Hoffmann, Adrian, Birk Diedenhofen, Bruno Verschuere, and Jochen Musch. 2015. "A Strong Validation of the Crosswise Model Using Experimentally-Induced Cheating Behavior." *Experimental Psychology* 62:403-14.
- Hoffmann, Adrian, and Jochen Musch. 2015. "Assessing the Validity of Two Indirect Questioning Techniques: A Stochastic Lie Detector versus the Crosswise Model." *Behavior Research Methods* (online first).
- Höglinger, Marc, Ben Jann, and Andreas Diekmann. 2014. "Sensitive Questions in Online Surveys: An Experimental Evaluation of the Randomized Response Technique and the Crosswise Model." University of Bern Social Sciences Working Papers No. 9 (available from <https://ideas.repec.org/p/bss/wpaper/9.html>).
- Holbrook, Allyson L., and Jon A. Krosnick. 2010. "Measuring Voter Turnout by Using the Randomized Response Technique: Evidence Calling into Question the Method's Validity." *Public Opinion Quarterly* 74:328-43.

- Horton, John, David Rand, and Richard Zeckhauser. 2011. "The Online Laboratory: Conducting Experiments in a Real Labor Market." *Experimental Economics* 14:399-425.
- Ipeirotis, Panagiotis G. 2010. "Analyzing the Amazon Mechanical Turk Marketplace." *XRDS Crossroads* 17(2):16-21.
- Jann, Ben, Julia Jerke, and Ivar Krumpal. 2012. "Asking Sensitive Questions Using the Crosswise Model. An Experimental Survey Measuring Plagiarism." *Public Opinion Quarterly* 76:32-49.
- John, Leslie K., George Loewenstein, Alessandro Acquisti, and Joachim Vosgerau. 2013. "Paradoxical Effects of Randomized Response Techniques." Unpublished manuscript (available from <https://vlab.decisionsciences.columbia.edu/newsletter/Feb1113/Paradoxical%20Effects%20of%20Randomized%20Response%20Techniques.pdf>).
- Kirchner, Antje. 2015. "Validating Sensitive Questions: A Comparison of Survey and Register Data." *Journal of Official Statistics* 31:31-59.
- Korndörfer, Martin, Ivar Krumpal, and Stefan C. Schmukle. 2014. "Measuring and Explaining Tax Evasion: Improving Self-reports Using the Crosswise Model." *Journal of Economic Psychology* 45:18-32.
- Kreuter, Frauke, Stanley Presser, and Roger Tourangeau. 2008. "Social Desirability Bias in CATI, IVR, and Web Surveys." *Public Opinion Quarterly* 72:847-65.
- Krumpal, Ivar, and Anatol-Fiete Näher. 2012. "Entstehungsbedingungen sozial erwünschten Antwortverhaltens." *Soziale Welt* 63:65-89.
- Kundt, Thorben C. 2014. "Applying 'Benford's Law' to the Crosswise Model: Findings from an Online Survey on Tax Evasion." Helmut Schmidt University Hamburg Department

- of Economics Working Paper No. 148 (available from http://www.hsu-hh.de/fgvwl/index_N0SfvfXiS58YR65c.html).
- Kundt, Thorben C., Florian Misch, and Birger Nerré. 2014. "Re-Assessing the Merits of Measuring Tax Evasions through Surveys: Evidence from Serbian Firms." ZEW Discussion Paper No. 13-047 (available from <http://ftp.zew.de/pub/zew-docs/dp/dp13047.pdf>).
- Lelkes, Yphtach, Jon A. Krosnick, David M. Marx, Charles M. Judd, and Bernadette Park. 2012. "Complete Anonymity Compromises the Accuracy of Self-Reports." *Journal of Experimental Social Psychology* 48:1291-99.
- Lensvelt-Mulders, Gerty J. L. M., Joop J. Hox, Peter G. M. van der Heijden, and Cora J. M. Maas. 2005. "Meta-Analysis of Randomized Response Research: Thirty-Five Years of Validation." *Sociological Methods & Research* 33:319-48.
- Mason, Winter, and Siddharth Suri. 2012. "Conducting Behavioral Research on Amazon's Mechanical Turk." *Behavior Research Methods* 44:1-23.
- Mood, Alexander M., Franklin A. Graybill, and Duane C. Boes. 1974. *Introduction to the Theory of Statistics*. New York: McGraw-Hill.
- Moshagen, Morten, Benjamin E. Hilbig, Edgar Erdfelder, and Annie Moritz. 2014. "An Experimental Validation Method for Questioning Techniques That Assess Sensitive Issues." *Experimental Psychology* 61:48-54.
- Moshagen, Morten, Jochen Musch, and Edgar Erdfelder. 2012. "A Stochastic Lie Detector." *Behavior Research Methods* 44:222-231.
- Ostapczuk, Martin, and Jochen Musch. 2011. "Estimating the Prevalence of Negative Attitudes towards People with Disability: A Comparison of Direct Questioning, Projective Questioning and Randomised Response." *Disability and Rehabilitation* 33:399-411.

- Peeters, Carel F. W. 2005. *Measuring Politically Sensitive Behavior. Using Probability Theory in the Form of Randomized Response to Estimate Prevalence and Incidence of Misbehavior in the Public Sphere: A Test on Integrity Violations*. Dissertation, Faculty of Social Sciences, Vrije Universiteit Amsterdam.
- Rosenfeld, Bryn, Kosuke Imai, and Jacob N. Shapiro. 2015. "An Empirical Validation Study of Popular Survey Methodologies for Sensitive Questions." *American Journal of Political Science* (online first).
- Shamsipour, Mansour, Masoud Yunesian, Akbar Fotouhi, Ben Jann, Afarin Rahimi-Movaghar, Fariba Asghari, and Ali Asghar Akhlaghi. 2014. "Estimating the Prevalence of Illicit Drug Use among Students Using the Crosswise Model." *Substance Use & Misuse* 49:1303-10.
- Smith, Tom W. 1992. "Discrepancies between Men and Women in Reporting Number of Sexual Partners: A Summary from Four Countries." *Social Biology* 39:203-211.
- Suri, Siddhartha, Daniel G. Goldstein, and Winter A. Mason. 2011. "Honesty in an Online Labor Market." Paper presented at the Twenty-Fifth AAAI Conference on Artificial Intelligence (available from <http://www.aaai.org/ocs/index.php/WS/AAAIW11/paper/view/3955/4262>).
- Tourangeau, Roger, and Tom W. Smith. 1996. "Asking Sensitive Questions: The Impact of Data Collection Mode, Question Format and Question Context." *Public Opinion Quarterly* 60:275-304.
- Tourangeau, Roger, and Ting Yan. 2007. "Sensitive Questions in Surveys." *Psychological Bulletin* 133:859-83.
- Warner, Stanley L. 1965. "Randomized-Response: A Survey Technique for Eliminating Evasive Answer Bias." *Journal of the American Statistical Association* 60(309):63-69.

Wolter, Felix, and Peter Preisendörfer. 2013. "Asking Sensitive Questions: An Evaluation of the Randomized Response Technique vs. Direct Questioning Using Individual Validation Data." *Sociological Methods & Research* 42:321-353.

Yu, Jun-Wu, Guo-Liang Tian, and Man-Lai Tang. 2008. "Two New Models for Survey Sampling with Sensitive Characteristic: Design and Analysis." *Metrika* 67:251-63.

Table 1. Descriptive statistics of the sample

		Percent
Gender	male	49.9
	female	50.1
Age	18 – 24	24.3
	25 – 29	27.0
	30 – 34	18.5
	35 – 39	10.7
	40 – 49	10.1
	50 or older	9.3
Education	college degree	54.0
	some college	34.2
	high school or other	11.8
Labor market status	employed	54.1
	self-employed	12.7
	unemployed	11.3
	student	13.0
	other	8.9
Current location	at home	85.4
	at work	9.9
	other	4.7
Prior MTurk studies	0	6.8
	1 – 9	19.3
	10 – 99	32.9
	100 – 999	30.2
	1000 or more	10.8

Notes: Labor market status recoded from multiple response data (prioritizing categories in the order as listed in the table); N = 6'152

Table 2. Sensitive questions

Item	Wording
Shoplifting	“Have you ever intentionally taken something from a store without paying for it?”
Tax evasion	“Have you ever provided misleading or incorrect information on your tax return?”
Non-voting *	“Did you vote in the 2012 US presidential election?”
Cheating in the prediction game *	“In the \$2 dice task at the beginning of this survey: Did you honestly report whether your prediction of the dice roll was right?”
Cheating in the roll-a-six game *	“In the \$2 dice game at the beginning of this survey: Did you honestly report whether you actually rolled a 6?”

* Reverse coded for the purpose of analysis.

Table 3. Number of observations by dice game variant and sensitive question technique

	Prediction game	Roll-a-six game
Direct questioning (DQ)	387	382
Crosswise-model RRT (CM)	1168	1145
Unrelated-question RRT (UQ)	760	780
Forced-response RRT (FR)	759	771

Table A1. Prevalence estimates by sensitive question technique as displayed in figure 1 (standard errors in parentheses)

	Shoplifting (N = 6136)	Tax evasion (N = 6136)	Non-voting (N = 6131)	Cheating in the prediction game (N = 3065)	Cheating in the roll-a-six game (N = 3070)
Direct questioning (DQ)	40.23 (1.77)	10.03 (1.08)	34.46 (1.72)	2.33 (0.77)	3.94 (1.00)
Crosswise-model RRT (CM)	46.42 (1.62)	19.52 (1.50)	38.11 (1.61)	15.41 (2.05)	14.34 (2.06)
Unrelated-question RRT (UQ)	54.53 (1.64)	17.63 (1.42)	34.74 (1.60)	3.74 (1.63)	5.23 (1.66)
Forced-response RRT (FR)	49.22 (1.71)	14.30 (1.52)	32.51 (1.68)	0.85 (1.83)	-1.94 (1.73)
Differences:					
CM – DQ	6.18 (2.40)	9.50 (1.85)	3.64 (2.35)	13.08 (2.19)	10.40 (2.29)
UQ – DQ	14.30 (2.41)	7.60 (1.78)	0.28 (2.34)	1.41 (1.80)	1.29 (1.94)
FR – DQ	8.99 (2.46)	4.27 (1.87)	-1.95 (2.40)	-1.47 (1.99)	-5.87 (2.00)

Table A2. Cheating rates in the prediction game and the roll-a-six game as displayed in figure 2 (standard errors in parentheses)

	Prediction game (N = 3065)			Roll-a-six game (N = 3070)		
	observed	estimated	difference	observed	estimated	difference
Direct questioning (DQ)	23.64 (2.50)	2.33 (0.77)	-21.32 (2.47)	4.46 (1.06)	3.94 (1.00)	-0.52 (0.74)
Crosswise-model RRT (CM)	26.63 (1.45)	15.41 (2.05)	-11.22 (2.42)	6.04 (0.71)	14.34 (2.06)	8.30 (2.08)
Unrelated-question RRT (UQ)	26.13 (1.80)	3.74 (1.63)	-22.40 (2.30)	5.01 (0.78)	5.23 (1.66)	0.21 (1.65)
Forced-response RRT (FR)	26.53 (1.80)	0.85 (1.83)	-25.68 (2.48)	5.20 (0.80)	-1.94 (1.73)	-7.14 (1.74)

Table A3. Individual-level validation results in the prediction game and the roll-a-six game as displayed in figure 3 (standard errors in parentheses)

	Prediction game (N = 3065)			Roll-a-six game (N = 3070)		
	TPR	FPR	CCR	TPR	FPR	CCR
Direct questioning (DQ)	9.84 (3.22)	0.00 (0.00)	78.68 (2.47)	70.59 (11.39)	0.82 (0.47)	97.90 (0.75)
Crosswise-model RRT (CM)	28.36 (5.52)	10.71 (2.64)	73.07 (2.93)	52.93 (9.46)	11.86 (2.08)	86.01 (2.05)
Unrelated-question RRT (UQ)	14.80 (4.70)	-0.18 (1.94)	77.74 (2.05)	54.77 (10.39)	2.61 (1.60)	95.25 (1.64)
Forced-response RRT (FR)	8.93 (5.05)	-2.07 (2.31)	75.84 (2.18)	41.11 (10.66)	-4.30 (1.69)	96.94 (0.73)

Notes: TPR = true positive rate, FPR = false positive rate, CCR = correct classification rate (negative false positive rates were set to zero for the computation of CCR)

Figure 1. Comparative validation of sensitive question techniques (point estimates and 95% confidence intervals)

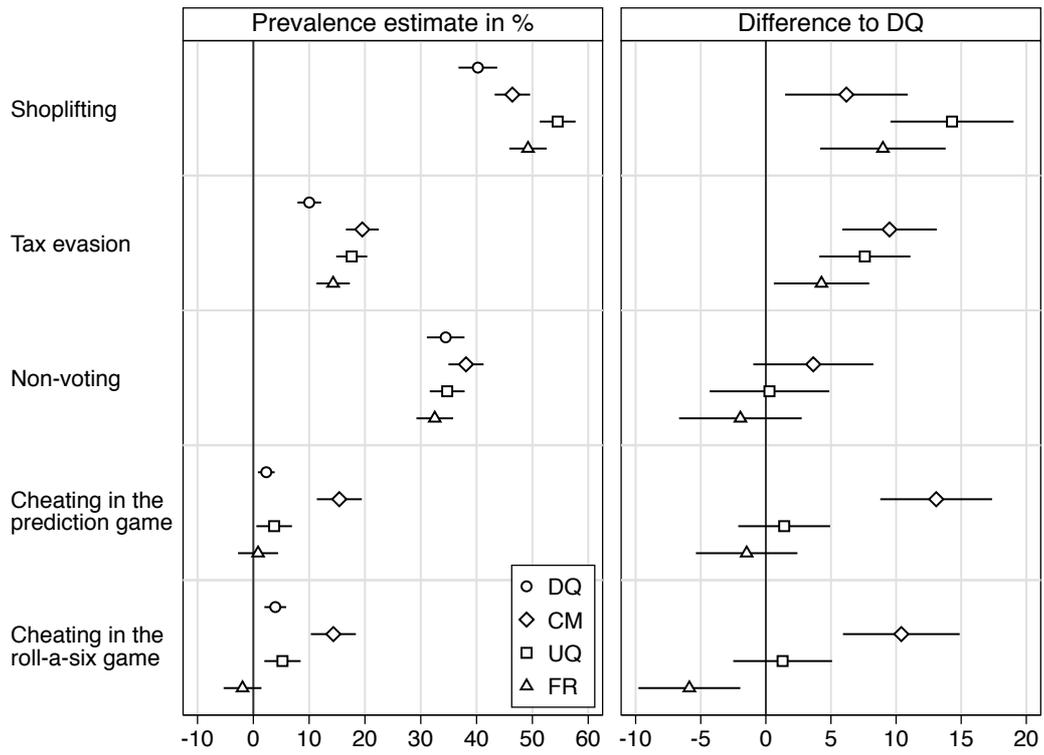


Figure 2. Aggregate-level validation of sensitive question techniques (point estimates and 95% confidence intervals)

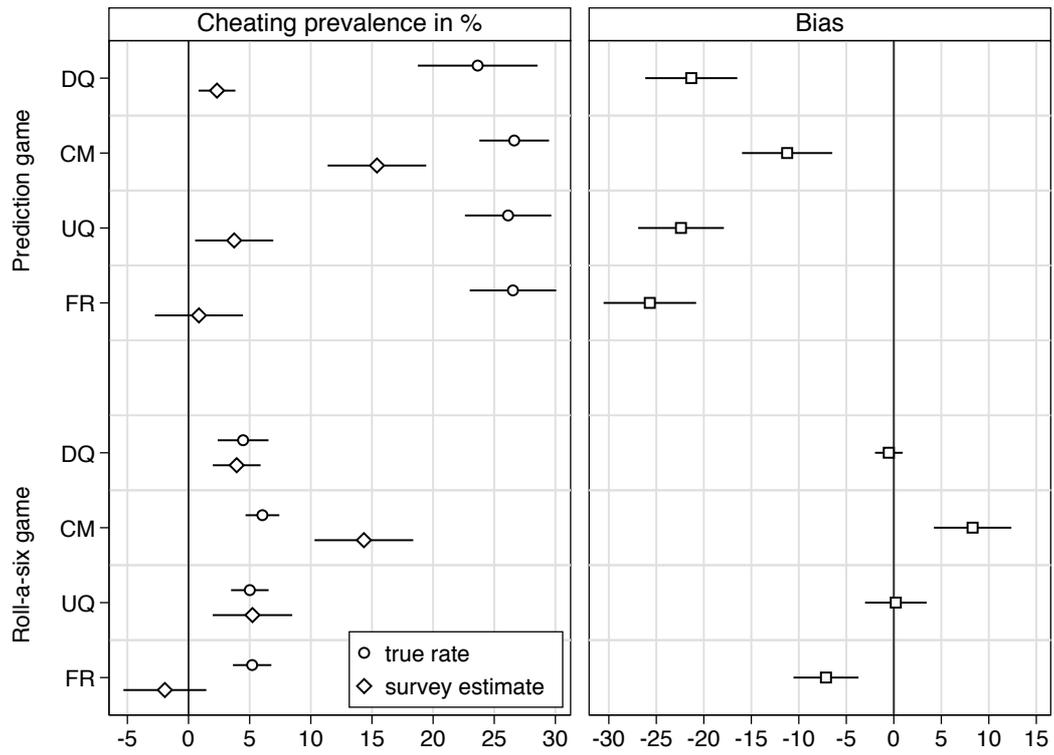
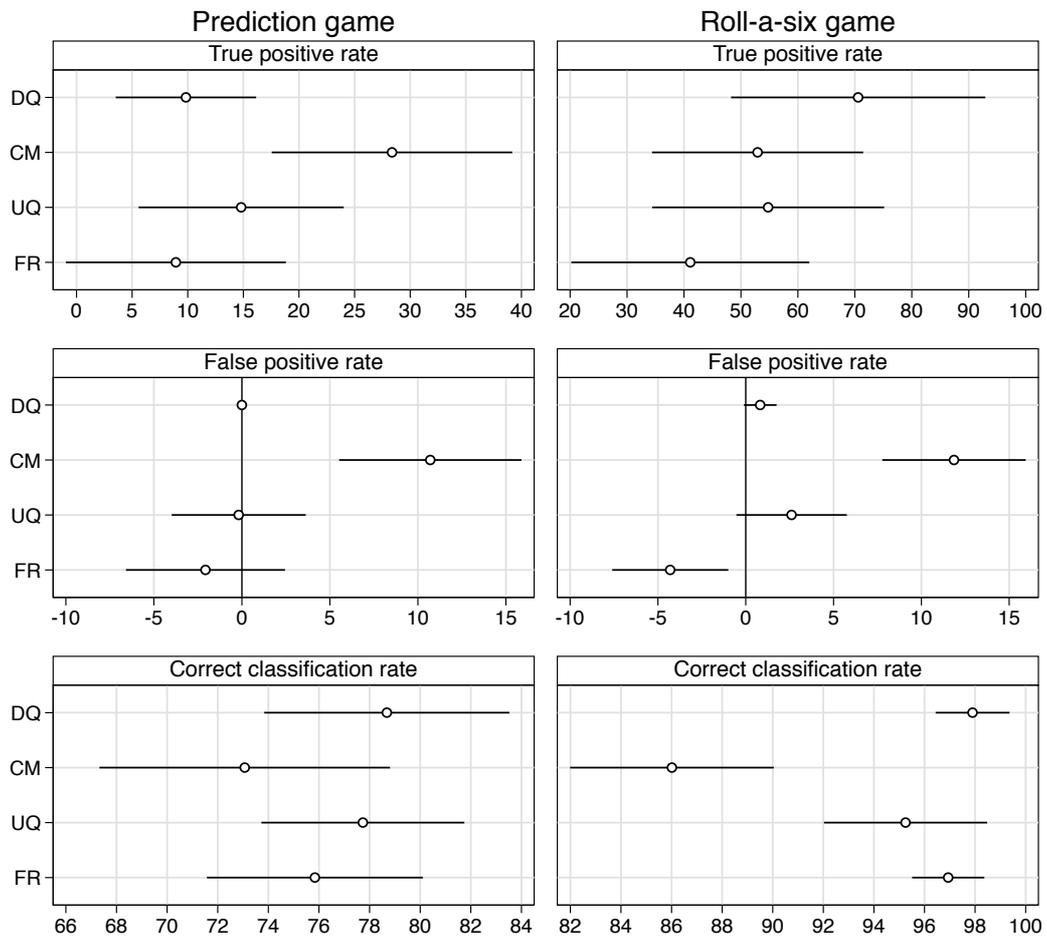


Figure 3. Individual-level validation of sensitive question techniques (point estimates and 95% confidence intervals)



Note: Negative false positive rates were set to zero for the computation of the correct classification rate.