

From the research bench

What is social desirability bias?

Social desirability bias pushes survey answers toward what looks good. It is systematic, so a larger sample keeps it rather than removing it.

Author	Published	Reading time	Words
Sinan Isoglu, MBA (Quantic) Commercial growth leader, lecturer and doctoral researcher	18 October 2026	14 min	3,037

Management summary

Social desirability bias is the tendency of survey respondents to underreport what is frowned upon and overreport what is admired. This concept guide defines it (Krumpal; Krosnick), separates a sensitive question from a sensitive answer, shows from the German General Social Survey (ALLBUS) which questions people leave blank, explains why the bias does not average out, reviews which survey designs reduce it and where they fail, sets out the evidence that some apparent lying is forgetting, and gives a one-page plan for asking a sensitive question. The sources are two literature reviews, mostly of general-population surveys about voting, drugs, sex and crime.

Contents	§ 1	What is social desirability bias?	2
	§ 2	Which answers count as sensitive?	2
	Figure 1	Item nonresponse in the German General Social Survey, by question and wave	4
	§ 3	Why does it shift an estimate rather than add noise?	5
	§ 4	What makes a respondent shade an answer?	5
	§ 5	Which design choices reduce misreporting?	6
	§ 6	Is every overreport a lie?	8
	§ 7	What should a survey plan settle before a sensitive question is asked?	9
	Figure 2	A sensitive-question plan	9
	§ 8	What have these reviews not tested?	11
		References	12

What is social desirability bias?

§ 1 of 8

It is the gap between what respondents did or think and what they report, when the report leans toward what is approved of. Krumpal (2013)² states the core in one sentence: “Due to self-presentation concerns, survey respondents underreport socially undesirable activities and overreport socially desirable ones.” Krosnick (1999)¹ describes the same pattern: “One well-known phenomenon in survey research is overreporting of admirable attitudes and behaviors and underreporting those that are not socially respected.”

¹ Krosnick, J. A. (1999). Survey research. *Annual Review of Psychology*, 50, 537–567.

² Krumpal, I. (2013). Determinants of social desirability bias in sensitive surveys: A literature review. *Quality & Quantity*, 47(4), 2025–2047.

The bias is not only a trait of some people. Krumpal reports a distinction drawn by Randall and Fernandes (1991): social desirability as a stable personality characteristic, “a constant need for social approval” that some respondents carry into every survey, and social desirability as an item characteristic, where “effects of social desirability are strongly influenced by characteristics of a specific item”. He adds that “aspects of the survey design, the interviewer’s characteristics and the survey situation determine the occurrence and the degree of social desirability bias”. That is the useful part for anyone who runs a survey: part of the bias is set by choices the researcher makes.

In this guide, social desirability bias means a systematic difference between true values and survey answers, in the direction the respondent expects to be judged more favourably, caused by how the answer would look to the interviewer, to others or to the respondents themselves. It is one kind of measurement error, and it is distinct from refusing to answer, from forgetting and from agreeing with whatever is asked.

Which answers count as sensitive?

§ 2 of 8


Not every awkward question. Krumpal reports a three-part distinction from Tourangeau and Yan (2007). The first is intrusiveness: “Questions asking about the respondents’ sexual preferences, health status or income are often perceived as too intrusive.” The second is threat of disclosure, the fear of what happens if the answer reaches others: “Such negative consequences could be: job loss, family upset or even prosecution.” The third is social desirability itself.

The third differs from the first two in one respect that matters for design. For social desirability, the problem “is not the sensitivity of a question but the sensitivity of an answer”. A question about drunk driving is easy for someone who never does it and hard for someone who does, so the bias falls on exactly the people the question is meant to count. Norms also differ between groups: Krumpal notes that “the tendency to give socially desirable responses may vary across cultural orientations”.

One indirect sign of sensitivity is a blank. Krumpal reports item nonresponse, the share of respondents who give no answer to a question, for seven questions in three waves of the German General Social Survey (ALLBUS), a national random sample of German-speaking adults in private households in Germany. Assuming that more sensitive items are left blank more often, he reads the table this way: “some items (household net income and voting intention) have consistently more missing data than other items”, and “the income question has the highest sensitivity of all items with nonresponse rates ranging from 20.7 to 26.2%”. Age and employment status were left blank by less than half a per cent. Figure 1 plots the table.

Figure 1 Item nonresponse in the German General Social Survey, by question and wave

About a fifth to a quarter of respondents left the income question blank in each wave. Age and employment status were left blank by less than half a per cent.



A blank is a sign of sensitivity, not a measure of misreporting: those who answer may still shade their answer. Voting intention was 14.4, 22.7 and 14.2 per cent; the five lower questions stayed below 2 per cent in every wave, so their dots overlap.

Source Krumpal (2013), p. 2027, Table 1, reporting item nonresponse in ALLBUS 1990, 2000 and 2006. The voting intention figures include the answer 'don't know'. The modes differ between waves: paper interviews in 1990, paper and computer-assisted interviews in 2000, computer-assisted personal and self-interviews in 2006.

The table needs two cautions. The voting figures carry a footnote: “Statistic includes answer category ‘don’t know’”. And the waves used different interview modes, so the change from one wave to the next is not a clean trend. Krumpal reports a second sign, ratings from Coutts and Jann (2011): in an online survey of 2,075 members of a German access panel, they asked whether each of several behaviours was all right and whether admitting it would be uncomfortable for most people, and added the two shares into one sensitivity score. “The most sensitive topics turned out to be shoplifting (79%) and infidelity (73%)”, against scores of 22 per cent for freeriding (travelling without a ticket) and 20 per cent for keeping too much change. Because each score adds two shares, it is not the share of respondents who found a topic sensitive.

Why does it shift an estimate rather than add noise?

§ 3 of 8

Because it has a direction. Krumpal, following Tourangeau et al. (2000), writes an answer as the true value plus a bias term plus random error, and draws the consequence: “Whereas random error cancels out over repeated measurements, response bias does not. Rather, the systematic difference between observed scores and true scores persists.” On this guide’s reading, the same holds across respondents: a survey four times as large halves the margin of error and leaves the bias where it was, so it estimates the wrong value more precisely.

The direction depends on what is approved of. Krumpal lists underreporting of drug use, smoking, alcohol and abortion, and “Respondents also underreport crime victimization, certain types of income (e.g. welfare) and unpopular attitudes, like racism and anti-Semitism.” Overreporting appears for voting, seat belt use, recycling and energy saving, and religious participation (p. 2030). At the level of the whole sample, “In the case of socially undesirable activities, sample proportions will underestimate the true prevalence and frequency of the frowned upon activities in question”, and “In the case of socially desirable behavior, an overestimation of the true level will occur.”

Krosnick’s¹ standard example is turnout: “the percentage of survey respondents who say they voted in the last election is usually greater than the percentage of the population that actually voted”. The [measurement error guide](#) separates an indicator from the construct it stands for; a self-report shaded toward approval is one way the two drift apart.

What makes a respondent shade an answer?

§ 4 of 8

Krumpal reviews two psychological mechanisms. In impression management, “respondents strive for social approval via selecting the answer that is expected to maximize positive valuations and minimize negative reactions by other subjects”. In self-deception, the audience is internal: “respondents themselves are the main addressees of socially motivated misreporting”. On this guide’s reading, the first responds to privacy and the second may not, because it travels with the respondent into any mode.

A rational-choice reading, also reviewed by Krumpal, treats the answer as a decision weighed on expected costs. It names three conditions that must all hold: a strong desire for approval, some perceived chance

of sanctions because the setting is not private, and a belief that the answer will change how others react (pp. 2030 to 2031). Vignette studies support the cost view, and the practical conclusion is plain: “As a baseline, lack of privacy lowers the respondent’s willingness to self-report norm-violating behavior.” Krumpal’s² summary puts the determinants together: “The degree of social desirability bias in sensitive surveys depends on the perceived items’ sensitivity, the degree of privacy in the interview situation, the fraction of the population who have behaved in the socially undesirable manner, and on specific aspects of the survey design.”

Which design choices reduce misreporting?

§ 5 of 8

Most of them raise privacy, and each has a known weakness.

Who asks. Self-administration is a well-supported lever: in many of the studies Krumpal reviews, self-administered modes raised reports of stigmatised conditions and behaviours, and in his summary, methods “minimizing the presence of the interviewer, seem to increase respondents’ privacy” and so produce more honest answers to sensitive questions. Where an interviewer is needed, the sealed envelope technique hands the sensitive part over on paper, so that “the interviewers do not have to read the embarrassing question and also remain ignorant about the respondents’ answers”. It is not complete: “In spite of the increased privacy, some respondents still refuse to fill out the confidential questionnaire”. A self-administered survey removes the interviewer, not every reason for the bias.

Who else is present. For bystanders, a meta-analysis that Krumpal reports found the effect of a spouse being present not significant, but “In contrast, the presence of parents generated more socially desirable answers.” In that meta-analysis of nine surveys, none of them an experiment, who the bystander was mattered more than whether there was one.

How it is asked and framed. Advice on forgiving wording rests mostly on survey practice, and the one recent experiment Krumpal reports shows inconsistent results. Confidentiality assurances help a little: a review found better reporting of sensitive items, “although the average effect size was small”, and Krumpal concludes that “Overall, data protection assurances seem to reduce respondents’ concerns and to improve response quality.” Some authors argue that elaborate assurances “might have unintended effects in terms of heightening respondents’ suspicions and concerns about who might get access to the data”; the experiments he cites for that concerned nonsensitive

topics and nonresponse. The bogus pipeline, in which “respondents believe that an objective procedure (e.g. lie detector, biochemical test) will be used to reveal false self-reports independent of whether such verification actually takes place or not”, relies on a deception. “The empirical evidence in many cases supports the assumed effectiveness of bogus pipeline, although some studies show unexpected results”, and his summary calls its effectiveness “still controversial”.

Indirect questions. The randomized response technique lets a random device decide which of two questions the respondent answers, and, in Krumpal’s² words, “Since the interviewer is unaware of the outcome of the random experiment, a given answer does not reveal anything definite about the respondent’s true status.” A meta-analysis reported by Krumpal concludes that “self-reports on sensitive issues are more accurate and more socially undesirable answers are elicited when RRT is employed”. Yet “standard direct questioning sometimes elicited more socially undesirable answers than did the RRT”, and “a substantial proportion of respondents do not comply with the RRT instructions”. The unmatched count technique hides the sensitive item in a list and asks only for a count; some studies found more employee theft and shoplifting with it than with direct questions, but for cocaine use studies found “significantly lower survey estimates for the UCT compared to standard direct questioning”, which Krumpal suggests could have come from the counting itself, quoting the original authors: “Deciding which items apply and keeping a running tally of these as the list is read has proven difficult for some respondents.”

How does a study know which method worked? Ideally, a validation study compares each answer with a record. Krumpal writes that these can be considered “the golden standard for testing the value of a method”. Without records, studies fall back on an assumption: “When underreporting is expected, a higher survey estimate (mean or proportion) is assumed to be a more valid estimate for the population parameter.” and “When overreporting is plausibly assumed, a lower estimate is interpreted as more valid.” On this guide’s reading, that assumption is the weak point of comparative studies: they show which method moved the estimate, not which estimate is true. Krumpal’s² verdict has two halves: “Improving the validity of answers to sensitive questions has proved a difficult task”, yet “Social desirability bias could be reduced by appropriately tailoring the survey design.”

Is every overreport a lie?

§ 6 of 8

No, and Krosnick (1999)¹ makes the case with the standard example. Two turnout patterns are usually read as lying: more respondents say they voted than the population's turnout allows, and "claims by significant numbers of people that they voted are not corroborated by official records". "In fact, these two empirical patterns are not fully attributable to intentional misrepresentation." Part of the gap comes from how turnout is calculated and from errors in the official records, part from who agrees to be interviewed, and part from pre-election interviews that make the people interviewed more likely to vote. "But even after controlling for all these factors, some people still claim to have voted when they did not."

What remains may not be lying either: "Surprisingly, recent research suggests that the widely believed explanation for this fact may be wrong." "Attempts to make people comfortable admitting that they did not vote have been unsuccessful in reducing overreporting", overreporters tend to be people who usually vote, and "the accuracy of turnout reports decreases as time passes between an election and a postelection interview". When an unpublished study warned respondents about memory confusion, it reduced overreporting. Krosnick's¹ conclusion is that "what researchers have assumed is intentional misrepresentation by respondents may be at least partly attributable instead to accidental mistakes in recall", and that "social desirability bias in questionnaire measurement may be less prevalent than has been assumed". This is a 1999 argument about one behaviour, and two of its steps rest on unpublished manuscripts; it shows that a memory error can pass for social desirability bias, not how much of a given gap each explains.

On this guide's reading, the remedy follows the cause: privacy does nothing for a memory error, and a recall aid does nothing for a respondent who wants to look good. Krosnick names a third look-alike, acquiescence, "the tendency to endorse any assertion made in a question, regardless of its content", which he puts at "an average acquiescence effect of about 10%" across the methods he reviews. One explanation he reports, from sociologists, is deference: "When researchers and interviewers are perceived as being of higher social status, respondents may defer to them out of courtesy and respect". The evidence on status is split, though, and the explanation he finds best supported is satisficing: agreeing because the respondent does not work the statement through. Either way, an agree/disagree item

can collect agreement for reasons that have nothing to do with approval of the answer.

What should a survey plan settle before a sensitive question is asked?

§ 7 of 8

Six things, in one row per sensitive item: which answer would be uncomfortable to give, the direction of the expected error, who would see the answer, who else could be present, the asking technique, and the check that will tell whether it worked. The fields follow the determinants and methods Krumpal reviews and Krosnick's¹ warning about other causes.

Figure 2 A sensitive-question plan

Write down the answer at risk, the expected direction, the audience, the setting, the technique and the check before the question goes out.

Answer	Direction	Audience	Setting	Technique	Check
Which answer would be uncomfortable to give?	Under- or overreport expected?	Who sees the answer with a name on it?	Who else could be present?	Direct, self-completed, list or random?	A record to compare, or only an assumption?

If the error could be memory rather than motive, plan a recall aid, not more privacy.

One row per sensitive item. A plan with no check can only assume that a higher or lower number is the truer one.

Source Author's decision worksheet. The answer and direction fields follow Krumpal (2013), pp. 2028 to 2030; the audience, setting and technique fields follow Krumpal (2013), pp. 2034 to 2040; the check field follows Krumpal (2013), p. 2033, and Krosnick (1999), pp. 545 and 546. The fields are reader inputs; no company data is supplied.

In customer research the sensitive answer is often not about drugs or income. On this guide's reading, it is the answer that would reflect badly on the respondent in front of the vendor or the employer: not

using a feature the company paid for, working around a required process, or having evaluated a competitor. The asker is often an account manager, whom the respondent will meet again.

An illustration, with no company data: a B2B software firm wants to know how many users at its customers log every customer call in the CRM, as their own sales process requires. The plan expects overreporting, because logging is the approved behaviour and the question is easy to answer with a yes. It moves the question from the account manager's review call to a web form that asks for no name, keeps team leads out of the distribution, and asks how many customer calls the respondent had last week and how many of them were logged, to limit recall error. Its check links no answer to a person: for a few accounts that agree, each large enough that no answer can be traced back, it compares the share of respondents who report logging all of last week's calls with the share of users whose calls that week were all logged, counted by the account from its own records. Those records have to include the calls made, from a phone or calendar system, because a call that was never logged leaves no trace in the CRM. And because the records cover every user while the survey covers only those who answered, part of any gap can come from who answered. Linking named answers to activity logs would undo the privacy the form promises and turn the survey into a check on individual employees. The account's own count also looks at each user's records, so even this design raises data-protection questions that this guide does not answer. Announcing any comparison may change the answers as well, which is what the bogus pipeline relies on. If the two shares agree, the firm learns that a nameless form works for this topic in those accounts; if not, it learns roughly by how much the answers miss there. These are the author's examples and have not been tested.

The [survey bias essay](#) files social desirability under wording, one of several routes from a sample to a wrong decision; this guide adds the levers beyond wording: who asks, who else is present and how the answer is collected.

What have these reviews not tested?

§ 8 of 8

Customer and employee surveys inside companies. Krumpal (2013)² reviews general-population, student and youth surveys on drugs, sex, crime, voting, attitudes and income; the business cases he cites are employee theft and premature sign-offs on audits. Krosnick (1999)¹ reviews survey methods in general, and his evidence on social desirability concerns turnout. Both are reviews, not new studies, and neither gives a correction factor that could be applied to a result after the fact.

On this guide's reading, two things carry over to B2B research with some confidence: the direction of the bias follows what the respondent believes the asker, or anyone who could see the answer, approves of, and privacy from them matters. How large the effect is in a customer survey, and which technique works there, the sources do not say.

 End of essay

References

- 1 Krosnick, J. A. (1999). Survey research. Annual Review of Psychology, 50, 537-567. <https://doi.org/10.1146/annurev.psych.50.1.537>
- 2 Krumpal, I. (2013). Determinants of social desirability bias in sensitive surveys: A literature review. Quality & Quantity, 47(4), 2025-2047. <https://doi.org/10.1007/s11135-011-9640-9>

How to cite this essay

Isoglu, S. (2026, October 18). What is social desirability bias? Sinan Isoglu. <https://isoglu.com/insights/journal/what-is-social-desirability-bias/>
Online version: <https://isoglu.com/insights/journal/what-is-social-desirability-bias/>

Sinan Isoglu, MBA (Quantic)

Commercial growth leader, lecturer and doctoral researcher
© 2026 Sinan Isoglu · isoglu.com · All rights reserved