

ORIGINAL ARTICLE

Who cares? Measuring differences in preference intensity

Charlotte Cavaillé¹ , Daniel L. Chen² and Karine Van der Straeten²

¹Gerald R. Ford School of Public Policy, University of Michigan, Ann Arbor, USA and ²Department of Social and Behavioral Sciences, Toulouse School of Economics (TSE), CNRS, Institute for Advanced Study in Toulouse (IAST), Toulouse, France
Corresponding author: Charlotte Cavaillé; Email: cavaillé@umich.edu

(Received 10 April 2023; revised 16 October 2023; accepted 8 December 2023)

Abstract

How well do existing survey instruments differentiate between opinions that affect individual behavior and opinions that don't? To answer this question, we randomly assigned U.S. respondents to one of three survey instruments: Likert items (Likert), Likert items followed by personal importance items (Likert+) and Quadratic Voting for Survey Research (QVSR), which gives respondents a fixed budget to buy "favor" or "oppose" votes, with the price for each vote increasing quadratically. We find that, relative to Likert, both Likert+ and QVSR better identify people who care enough about an issue to act in opinion-congruent ways, with QVSR offering the most consistent improvement overall. Building on these results, we show how conclusions regarding the relationship between policy opinions and self-interest can differ across measurement strategies.

Keywords: Attitude extremity; attitude importance; determinants of policy preferences; likert item; survey research, preference intensity; public opinion; quadratic voting for survey research; self-interest

People have opinions on many issues but do not care about all of them. If people care deeply about an issue, they are more likely to act in opinion-congruent ways, e.g., object (through voting or protest) when a proposed policy does not align with what they prefer. If people are mostly indifferent, then they are more likely to compromise, e.g., accept an outcome that does not align with what they prefer (by going to the beach instead of protesting). Conceptually, these differences across opinions—which we will later call differences in preference intensity—are central to theories of democratic accountability (Hill, 2022), issue voting (Rabinowitz and Macdonald, 1989) and party-switching (Carsey and Layman, 2006).

How do researchers empirically differentiate between opinions that affect political behavior and opinions that don't? One type of measurement strategy sequentially asks people how they feel about a set of issues, with responses recorded using a uni- or bi-directional scale. This simply-ask approach assumes that respondents, when prompted with words such as "strongly" or "important," report to the best of their ability how much they care about an issue. An alternative approach seeks to elicit the same information by forcing respondents to trade-off across all issues considered jointly. In this case, expressing their opinion on some issues comes at the expense of doing so on others: respondent have to choose. Such forced-choice approach assumes that respondents arbitrate in ways that are informative of how people behave when confronted with the real world costs of opinion-congruent action.

From the existing research, we already know that the simply-ask approach helps distinguish people who care about an issue from those who don't: survey respondents who indicate feeling "strongly" about a policy they find "very important" are more likely to behave in

opinion-congruent ways than people who pick the “weakly” and “not important” response categories (e.g., Krosnick and Petty, 1995; Carsey and Layman, 2006). For people with well-formed opinions on issues asked about in the survey and no obvious reasons to misrepresent their “true” opinion, this approach should suffice. Yet, concerns that few respondents match this profile have led many scholars to distrust subjective survey data (Bertrand and Mullainathan, 2003). Forced-choice elicitation strategies could offer a compromise solution (Cavaillé *et al.*, 2019; Hanretty *et al.*, 2020). As we discuss later in the paper, they can help people better realize “in the moment” how important an issue is to them or mitigate the measurement bias introduced by competing motives such as partisanship. Alternatively, if concerns regarding the simply-ask approach are overblown, then applied researchers need not rush to look for better ways of identifying “who cares.”

To investigate these issues, we asked respondents from a representative sample of U.S. citizens their opinion on 10 policy issues, randomly varying the method used to measure their opinion. One method is the Likert item (Likert for short), which asks people to report (on a 3-point scale) the “strength” of their support/opposition. The other method (Likert+) combines the Likert item with a personal importance item that asks respondents to further specify if the issue previously mentioned is “personally important” to them (on a 5-point scale). The third method—Quadratic Voting for Survey Research (QVSR)—uses a variant of the forced-choice approach. It gives respondents a fixed budget to “buy” votes in favor or against each of the 10 policy proposals, with the price for each vote increasing quadratically. Because of this price schedule, it becomes increasingly costly to acquire additional votes to express more intense support or opposition to a given policy (Lalley and Weyl, 2018). After expressing their opinion using one of these three measurement tools, respondents performed a number of choice tasks commonly associated with issue-specific political action (e.g., a donation to a non-profit advocating for gun control or letter writing to a senator about a minimum wage bill). We compare each tool’s ability to distinguish between respondents whose opinion-congruent behavior suggests they care intensely about a given issue and those whose non-congruent behavior suggests they do not care as much.

First, we document Likert’s reasonably good performance. We find that the addition of a personal importance item offers some improvement, with QVSR offering the most consistent improvement overall. One important difference appears to be QVSR’s ability to de-bunch in informative ways, that is, generate meaningful differences in votes cast among people who, under alternative measures, would end up picking the same response category.

Because of these differences, the measurement strategy used to measure people’s opinions has implications for applied research. We demonstrate this point by revisiting the claim, common among public opinion scholars, that people’s policy opinions do not reflect their material self-interest (e.g., Sears and Funk, 1990). In line with previous studies, we find that support for a policy measured using a Likert scale conveys only limited information about a respondent’s position as a potential beneficiary of this policy. In contrast, QVSR votes help distinguish between respondents who would directly benefit from a policy and respondents who would not be affected. This suggests that conclusions regarding the importance of material self-interest can vary with the measurement strategy used to measure individual support for a given policy.

1. Measuring who cares? Conceptual and theoretical considerations

Consider a status quo changing policy (e.g., Brexit). Of all people who favor this policy, when given the opportunity (e.g., a referendum), only a subset will translate this support into opinion-congruent action (e.g., turn out to vote in favor). Formally, we capture these individual differences with a real number u_{ik} in the interval $[-1, 1]$, where the likelihood of individual i taking costly action in favor of the reform k and against the status quo increases as u_{ik} gets closer to 1. Conversely, the likelihood of taking costly action against the reform and in favor of the status quo increases as u_{ik} gets closer to -1 . The preference ranking captured by u_{ik} can be further decomposed into two terms. One, preference orientation, is an indicator variable that captures

whether the respondent prefers the reform over the status quo ($u_{ik} > 0$) or the status quo over the reform ($u_{ik} < 0$). The other, preference intensity is the *extent* to which the respondent prefers one over the other and is captured by the absolute value of u_{ik} ($|u_{ik}|$).¹

As a summary concept aimed at describing a complex subjective mental state, u_{ik} cannot be observed, it can only be imperfectly measured. To recover meaningful information about u_{ik} in general, and $|u_{ik}|$ in particular, researchers rely on two broad families of measurement strategies. Next, we discuss the pros and cons of each. For expository purposes, we build our discussion around the specific versions used in the empirical section of this study.

1.1 The simply-ask approach: Likert and personal importance items

When measuring u_{ik} , researchers who favor the simply-ask approach described in the introduction most often rely on the two-step version of the Likert item (Malhotra *et al.*, 2009). First, respondents are asked if they “favor, oppose, or neither favor nor oppose” a status quo changing policy k . Respondents who pick the favor or oppose option then see the following prompt: “Do you favor [oppose] that a great deal, moderately, or a little?” Respondents who initially select “neither nor” are not asked a follow-up question. Recorded responses range from -3 (strongly oppose) to $+3$ (strongly favor) and are centered around 0 (neither-nor). Once normalized, the resulting response variable \hat{u}_{ik}^L ranges from 0 to 1.

A common practice is to supplement information provided by Likert items using a follow-up personal importance item. This item asks respondents “how important” a given issue is to them “personally.” Respondents answer using a categorical scale ranging from “not at all important” (1) to “extremely important” (5) (Miller and Peterson, 2004; Howe and Krosnick, 2017).² A recurrent finding is that opinion-congruent behavior is higher among people who “strongly favor” a policy and among those who report finding the issue personally important to them (e.g., Krosnick and Petty, 1995). In other words, both Likert and personal importance items recover meaningful information about u_{ik} in general and $|u_{ik}|$ in particular. Combined with a Likert item, this suggest a second straightforward way of measuring u_{ik} , namely:

$$\hat{u}_{ik}^{L+} = F\left(\hat{u}_{ik}^L, \widehat{Imp}_{ik}\right)$$

with the answers to the personal importance item denoted by \widehat{Imp}_{ik} .

Before discussing the pros and cons of this simply-ask approach, a quick note on how preference intensity, per our definition, relates to similar concepts in public opinion research. We have defined our main quantity of interest in reference to a spatial model of politics most commonly found in political economy (see Appendix B.1 for more details). Likert and personal importance items were developed by social psychologists to measure what is called “attitude extremity” and “attitude importance.” While attitude extremity captures “the degree to which the person likes or dislikes the object,” attitude importance captures “an individual’s subjective judgment of the significance he or she attaches to his or her attitude” (Howe and Krosnick, 2017, 329).³ What is the

¹Both preference orientation and intensity vary with the specific content of the proposed policy and how it compares to the status quo. Preference intensity, in addition, varies with the relative importance of the policy domain. One individual might feel strongly about the policy domain, yet be quite satisfied with the status quo, resulting in moderately intense preferences. Another might not care much about a policy domain, yet be dissatisfied enough with the status quo to be moderately supportive of the reform. Distinguishing between these two components is not the focus of our analysis, though see appendices B.1 and B.2 for a more detailed discussion of how these relate to Likert and personal importance items.

²Per this definition and measurement strategy, personal importance is distinct from “national importance” which captures respondents’ subjective evaluation of how important a policy is “for the country as a whole” (Miller *et al.*, 2017, 157).

³Krosnick and Abelson (1992) identify at least five attitude features. Given our interest in explaining “attitude-congruent” behavior, we leave aside attributes (e.g., attitude accessibility) most relevant for cognition and attitude formation only (Howe and Krosnick, 2017).

relation between u_{ik} (the combination of preference orientation and intensity) on the one hand, and attitude extremity and importance on the other? There are significant epistemological differences underpinning spatial models' emphasis on preference ranking and social psychology's emphasis on attitudes, which we discuss in Appendix A. Still, for our purpose, we can put these differences aside: u_{ik} is, by construct, the *total sum effect* of attitude extremity, attitude importance and any other attitude features that affect the decision to act in an opinion-congruent way or compromise instead.⁴ Our goal is to *measure* differences in u_{ik} to the best of our abilities, not to explain these differences. As a result, attitude extremity and importance are absent from our conceptualization or analysis: preference intensity supersedes these concepts. Note that a concept such as attitude strength, which Krosnick and Abelson (1992) define as the extent to which a given attitude "affects one's cognition or behavior," does not provide an adequate substitute for preference intensity as defined here. One important reason is that social psychologists relate strong attitudes to stable attitudes that are hard to change. In contrast, based on our definition of $|u_{ik}|$, preference intensity can vary over time depending, for example, on changes in the status quo.⁵

The main advantage of the simply-ask approach is its simplicity. One major disadvantage for researchers interested in measuring preference intensity $|u_{ik}|$ is that it puts respondents in a world where talk is cheap.⁶ First, there are no consequences for misrepresenting one's true opinion or reporting an opinion even if one has none. A second concern is that respondents are asked about policy issues sequentially, with no incentives to arbitrate between intense preferences for two mutually exclusive policies. If people have some prior sense of how their opinion on one issue compares to their opinion on another and report these truthfully, these concerns would be relatively minor. But scholars have reasons to worry. Partisan motives have been shown to systematically bias survey responses (Bullock and Lenz, 2019). In the U.S. context, polarized ideological messaging and affective partisanship can generate bi-modal response distributions. In this case, the same response category (e.g., "favor a great deal" or "very important") might include respondents who care about the issue and respondents who do not care as intensely and are merely "paying lip service to the party norm" (Zaller, 2012). Not only do researchers have limited variation to build on; whatever variation they have is difficult to interpret. Furthermore, with only two parties to choose from, many U.S. voters hold a combination of mutually exclusive policy preferences: behaving in an opinion-congruent way on one policy often means having to compromise on another (e.g., support for Republicans' strong stance on balanced budgets means compromising on support for abortion rights). With the simply-ask approach, respondents are in a world of abundance, where compromise is not needed, meaning that the information recovered might carry too little information about opinion-congruent behavior in the real world.

1.2 Forced-choice approach: quadratic voting for survey research

These concerns have led some scholars to turn away from subjective survey data and stated preferences and rely instead on in-survey behavioral outcomes in the form, for example, of a donation or a real effort task. While ideal for studies limited to one or two issues, in-survey behavioral

⁴Given the different epistemological starting points, building a 1-to-1 conceptual match between attitude extremity and importance on the one hand, and related components of preferences on the other, is far from straightforward. We discuss this in more detail in Appendix B.2.

⁵The concept of preference intensity, per our definition, is also different from that of "attitude intensity," which social psychologists define as "the strength of the emotional reaction provoked by the attitude object" (Krosnick *et al.*, 1993, 1132). In contrast, preference intensity is defined only in reference to the *net* utility of the proposed change relative to the status quo, which might or might not be driven by an emotional response to an object.

⁶Note that this is not a concern for people interested in measuring attitudes, which are conceptually different from preference ranking as defined in this paper.

proxies are difficult and/or costly to scale up to include a larger number of issues. An intermediate solution is to rely on stated preferences, but use a measurement strategy that leverages a forced-choice design that makes talk a little less cheap by confronting people with trade-offs.

QVSR, developed by Posner and Weyl (2018), is one such measurement strategy.⁷ Like Likert items, it asks respondents the extent to which they favor a given set of policies, but the technology used to measure people's answers is very different. Respondents express their preferences on a bundle of policies under the constraint of a fixed budget of credits with which to buy units of support (votes in favor) and units of opposition (votes against).⁸ A distinctive feature of QVSR is that the price schedule is quadratic: buying one vote for one proposal costs one credit; buying two units for the same proposal costs four credits; buying three units costs nine credits; and so on. In our own survey, respondents assigned to QVSR were given a budget of 100 credits to spend across ten different survey questions. Figure 1 shows what this survey looks like to respondents. Respondents can scroll down to report their preferences on all the issues examined in the survey. Remaining credits are displayed at the top of the screen. Respondents can go back to revise their answers until they are satisfied with how they have allocated their credits. The maximum that respondents can spend in favor or against any question is 10 units of support/opposition (which costs 100 credits) though doing so would mean not being able to express (however mild) support for or opposition to any of the other 9 issues. Respondents do not have to spend all of their 100 credits.⁹ Recorded responses range in theory from -10 to +10. Once normalized, the resulting response variable \hat{u}_{ik}^{QVSR} ranges from 0 to 1.

QVSR's forced-choice design compels individuals to compare across issues. This can improve the quality of responses in three ways. First, QVSR better approximates the real-world opportunity costs of opinion-congruent behavior. Second, it does not require people to have well-formed opinions: by forcing people to compare across issues, QVSR can induce people to themselves realize what it is they care about the most. Third, as discussed in Appendix B, when partisan concerns generate misreporting and end-of-scale bunching, QVSR forces people to de-bunch in ways that are informative of preference intensity. In the abundance world of Likert and personal importance items, people can inflate their reported preference intensity at no cost. The combination of a fixed budget and quadratic pricing makes this type of misreporting costly: expressing a strong preference (through multiple votes) for a policy one does not care about comes at the cost of doing so for a policy one truly cares about. Take, for example, a set of respondents who all report strongly supporting unrestricted abortion and finding this issue personally important to them. Some might provide these answers because they are sincerely reporting their true u_{ik} . Others might have a lower u_{ik} yet choose end-of-scale responses out of partisan concerns (e.g., strong support for abortion rights is what defines a strong Democrat). Assuming respondents compromise (in terms of the number of votes cast) on policies they do not sincerely care about, then we can plausibly expect QVSR to be more informative of differences in preference intensity.

Still, QVSR has several important drawbacks. One is that it requires higher cognitive engagement from survey respondent, something that might improve the quality of responses for some but decrease it for others. For example, some respondents might find the instrument

⁷Another type of forced-choice method asks respondents to choose between bundles of policies (Hainmueller *et al.*, 2014; Hanretty *et al.*, 2020; Sides *et al.*, 2022). Because these methods return *group-level* estimates of preferences (instead of measuring individual-level differences), they are outside our scope of inquiry (Abramson *et al.*, 2022; Ganter, 2023).

⁸A related method asks respondents to rank policies by order of importance. When choosing a forced-choice method to evaluate, we opted for QVSR for three reasons. First, it jointly measures preference orientation and preference intensity. In contrast, a ranking exercise first requires a battery of Likert items to capture preference orientation. Second, ranking exercises tend to be limited to five items while piloting with QVSR shows that respondents are comfortable with a larger set of items (Quarfoot *et al.*, 2017). Third, QVSR has the potential to generate more information in the form of a cardinal (instead of an ordinal) scale.

⁹Description of how respondents interact with this interface is not the focus of our study, see Quarfoot *et al.* (2017) for more information.

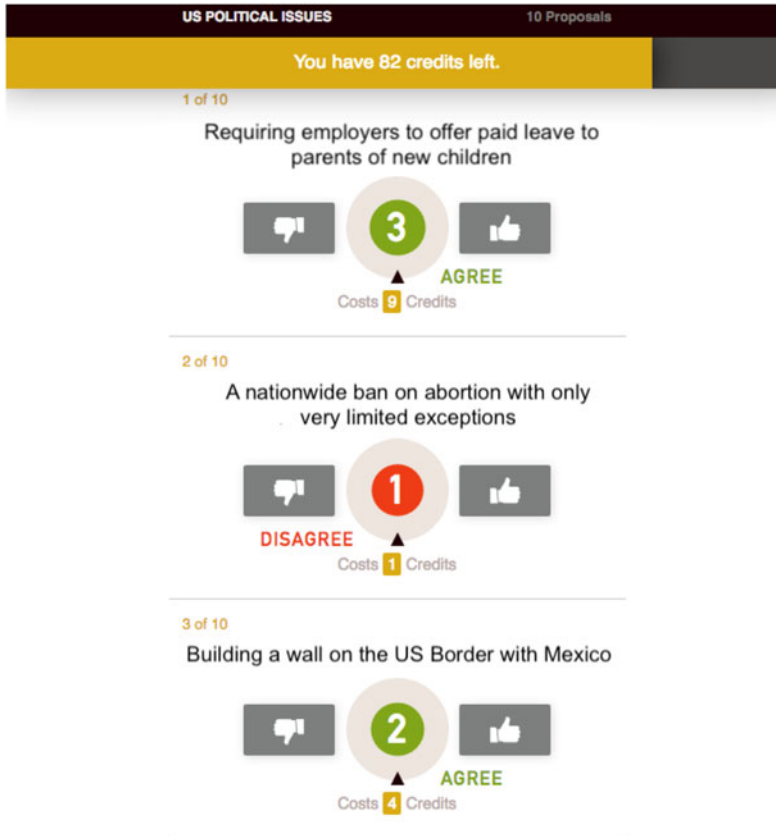


Figure 1. Screenshot of the QVSR Version of the Survey.

too demanding and respond using bias-inducing heuristics (Krosnick, 1991; Sauer *et al.*, 2011). For these respondents, a simpler simply-ask survey instrument would do a better job. A second drawback is that, while plausibly approximating the type of arbitrage most relevant to the measurement of preference intensity, QVSR's budget constraint might also introduce measurement error. For example, if the fixed budget is too constraining, then respondents can end up randomly picking which issue to give fewer votes to in order to free enough credits for other issues. A related concern is that of interpersonal comparisons. Take, for example, two respondents who both used 9 credits (3 votes) to express support for a given proposal: can we reasonably assume that they care about this proposal to the same extent? Note that this issue is a concern for most subjective measurement tools. For example, with personal importance items, not everyone imparts the same meaning to the "extremely important" response category.

1.3 Comparing methodologies

How much is gained by measuring preference intensity using QVSR instead of Likert items? Assuming QVSR offers an improvement over Likert items, how does this improvement compare to merely adding a follow-up personal importance item? We conclude this overview by providing speculative, if informed, answers to these questions.

Likert provides our benchmark: to be of some value, a measurement strategy should perform better than simply (and sequentially) asking people how strongly they favor a given set of status

quo changing policies. Compared with Likert, we expect QVSR to provide a better measure of preference intensity. That's because QVSR is less prone to end-of-scale bunching and forces respondents to engage in between-issue comparison. How do Likert and QVSR compare to Likert+?

To compute Likert+, we multiply answers from the Likert and personal importance items. The resulting scale ranges from -15 to $+15$ ("strongly oppose/favor" and "extremely important") and is centered around 0 (neither-nor). Once normalized, the response variable \hat{u}_{ik}^{L+} ranges from 0 to 1. In Appendix E, we discuss alternative ways of combining the information captured by these two survey items. Results remain unchanged. We focus on the multiplicative approach because it aligns with prior evidence of a positive interaction between Likert answers and personal importance answers, meaning that people who strongly favor a policy are more likely to behave in opinion-congruent ways, increasingly so if this issue is personally important to them (Carsey and Layman, 2006; Miller *et al.*, 2017).¹⁰

Mechanically, given the addition of an item, researchers have more variation to work with when using Likert+ than when using Likert. Because Likert+ generates novel information in the form of a new prompt about a different facet of preference intensity, we expect Likert+ to outperform Likert. As we discuss in Appendix B, whether QVSR outperforms Likert+ partly depends on the strength of the partisan motive and its impact on the prevalence of uninformative end-of-scale bunching. It also depends on the amount of error introduced by respondents' heterogeneous reaction to QVSR's budget constraint. When it comes to comparing QVSR and Likert+, we remain agnostic on which methodology will outperform the other.

In Appendix C, we also offer a systematic comparison of the three methods focusing on costs (software, survey time and data loss rate). The creation of several QVSR web applications¹¹ have brought software costs down to zero. While median time spent answering preference-related questions is shorter for respondents assigned to Likert, it is roughly the same for respondents assigned to Likert+ and QVSR. The main difference time-wise for QVSR is a 90 second video explaining how the tool works.¹² QVSR, in our study, has one additional extra cost, namely a higher data loss rate (10 percent versus 4 percent for Likert and Likert+) due to duplicates and a higher dropout rate (though see Quarfoot *et al.*, 2017 who find no such difference).

As the above discussion suggest, each methodology comes with advantages and disadvantages. Which method outperforms the others is an empirical question, one to which we now turn.

2. Empirical design

A measurement tool can be thought of as a classification instrument that distributes the surveyed population across a fixed number of response categories. Each tool differs in terms of the number of available response categories and the technology used to distribute people across categories. The tool that best measures preference intensity is the one that best classifies respondents from the most to least likely to behave in an opinion-congruent way. To compare each survey tool's classification abilities, we use an experimental design. In this section, we first describe this design. Next, we describe how we use the data collected to compare Likert, Likert+ and QVSR.

2.1 Survey design

We asked people to take the same survey, randomly varying the measurement tool used to measure policy opinions. The survey was administered to a general population of U.S. citizens over the age of 18 ($N = 3551$). The survey company, GfK-Ipsos, uses a probability-based web panel designed to be representative of the U.S. population. The main data collection effort took

¹⁰See Appendix E for evidence using our own data.

¹¹See footnote 23 on page 24.

¹²The video can be found at: https://www.youtube.com/watch?v=GrY_RzDsqLY.

place from October 5 to October 9, 2018. For an overview of the survey design, see Appendix C and H.

Respondents were randomly assigned to one of the three survey tools and asked to provide their opinion on the following 10 policy issues:¹³

Do you Favor or Oppose:

- [sameS] Giving same sex couples the legal right to adopt a child
- [gunC] Laws making it more difficult for people to buy a gun
- [wall] Building a wall on the U.S. Border with Mexico
- [paidL] Requiring employers to offer paid leave to parents of new children
- [affA] Preferential hiring and promotion of blacks to address past discrimination
- [equalP] Requiring employers to pay women and men the same amount for the same work
- [minW] Raising the minimum wage to \$ 15 an hour over the next 6 years
- [abort] A nationwide ban on abortion with only very limited exceptions
- [cap] A spending cap that prevents the federal government from spending more than it takes
- [env] The government regulating business to protect the environment

After expressing their opinion, respondents were given the opportunity to take action by donating lottery money to single-issue advocacy groups. First, respondents were told that, as participants to the survey, they had been automatically entered into a lottery with “a prize of \$100 for 40 randomly selected respondents (among 4000 or so).” They were then prompted to imagine that they were among the winners and asked whether they wanted to donate part of their lottery money to an advocacy group. They had a choice between four advocacy groups working in two issue areas: immigration and gun control. For each issue area, we chose organizations that fall on different sides of the political divide: for and against immigration, as well as for and against gun control. Respondents could choose not to donate or to donate to one, and one only, of the four advocacy groups. Whatever they did not donate, they could keep. Two weeks after the end of the survey, 40 randomly selected respondents received their prize money, which was disbursed by GfK-Ipsos.¹⁴

Four months later (between January 31 and February 18, 2019), we recontacted a random subset of respondents and asked them to answer the same 10 survey questions using the survey tool they were assigned to in the first wave (number of responses, $N = 1569$).¹⁵ We then collected information on two additional behavioral tasks.

First, we asked each respondents how they would behave in three dictator games: one involving a Republican, another a Democrat and a third an Independent (the order was randomized). Respondents had the option to donate anywhere between \$0 and \$100 of some lottery money (the set up was similar to the one in wave 1). After they made their decisions, respondents were asked again about their donation to the Independent. We explained that, in wave 1, this Independent had donated to the pro-immigration organization and to the anti-gun control organization.¹⁶ We asked respondents if they wanted to change the amount they had previously decided to donate to this individual. In other words, they had to choose between doing nothing, “punishing” the Independent (by decreasing the amount originally donated) or “rewarding” them (by increasing

¹³In each treatment, the order in which the 10 proposals are presented is fully randomized.

¹⁴Because of regulations preventing tax-exempt research funds from being used for political purposes —something we failed to realized at the design stage of the study—we ultimately made no donations and, several weeks after the end of the study, the lottery winners received the full \$ 100 amount, alongside an email explaining the reason why.

¹⁵Participation in wave 2 is not predicted by treatment condition and policy preferences in wave 1, nor by partisanship. See Appendix C for more on balance across treatment conditions in wave 1 and on wave 2 participation.

¹⁶In practice, this was impossible as respondents could only donate once. When disbursing the funds, we consequently used survey answers, not donation decisions, to identify Independents to disburse the funds to.

Table 1. Behavioral outcomes and relevant survey question

Variable	Description	Mean (Stand. dev.)	Survey question expected to correlate with behavior
Donation to gun-related advocacy group	Equal to the \$ amount donated multiplied by 1 if donated to pro gun control and -1 if donated to anti gun control advocacy group.	9.8 (33.5)	Laws making it more difficult for people to buy a gun
Donation to immigration-related advocacy group	Equal to the \$ amount donated multiplied by 1 if donated to pro immigration and -1 if donated to anti immigration advocacy group.	1.6 (28.7)	Wall on the border with Mexico
Punishment of Independent respondent (1)	Equal to the \$ amount <i>taken off</i> the amount previously donated to the Independent. If respondent <i>gave</i> additional \$ then amount multiplied by -1.	4.9 (13.5)	Laws making it more difficult for people to buy a gun/ Wall on the U.S. border with Mexico
Punishment of Independent respondent (2)	Equal to the \$ amount <i>taken off</i> the amount previously donated to the Independent as a <i>proportion</i> of the amount originally donated. If respondent <i>gave</i> additional \$ then amount multiplied by -1.	0.17 (0.38)	Laws making it more difficult for people to buy a gun/ Wall on the U.S. border with Mexico
Letter writing on the minimum wage bill	Equal to the length of text written (number of characters).*	76 (139)	Raising the minimum wage to \$15/h over the next 3 years (absolute values)
Letter writing on the abortion bill	Equal to the length of text written (number of characters).*	59 (90)	A nationwide ban on abortion with only very limited exceptions (absolute values)

*See Wasow (2023) for more information on how to use text as a behavioral measure.

the amount originally donated). Because few people in our survey (based on wave 1 results) are both pro-immigration *and* anti-gun control, most respondents faced a trade-off: rewarding this fellow survey participant meant condoning a position one is in agreement with while also condoning a position one is in disagreement with.

Second, respondents were also given the opportunity to write to their Senators about real bills that were moving through Congress at the time of the survey. One bill was about abortion and the other was about raising the minimum wage. We did not mention who the bill sponsors were, only the content of the bills. The texts provided by the respondents were then integrated into a letter, which was ultimately sent to the Senate committees in charge of reviewing the policy proposals (Adida *et al.*, 2018). Comments were anonymous. This task was designed to capture respondents’ willingness to spend time and effort promoting a political cause they agree with.

As we discuss in Appendix C, in the QVSR treatment condition, dropout rates are higher by 5 percentage points (6 percent versus 1 percent). We found no evidence that dropping out was predicted by observable covariates including partisanship and ideology. Table 1 provides an overview of the outcome variables derived from the three behavioral tasks and used in the remainder of the analysis. Throughout the paper, when we examine the relationship between survey answers and behavior, we only use answers collected in wave 1.¹⁷ Using data collected in wave 2 does not change the results (See Appendix D).

2.2 Estimation strategy

Each survey tool generates a response variable (\hat{u}_{ik}^L , \hat{u}_{ik}^{L+} or \hat{u}_{ik}^{QVSR}) that differs from the other two in terms of (1) the total number of ordinal categories and (2) the distribution of observations across these categories. Likert has 7 response categories ranging from -3 to +3 and Likert+

¹⁷The reader should keep in mind however that, in at least two cases (dictator game and letter writing tasks), the attitudinal data is analyzed alongside behavioral data collected four months apart.

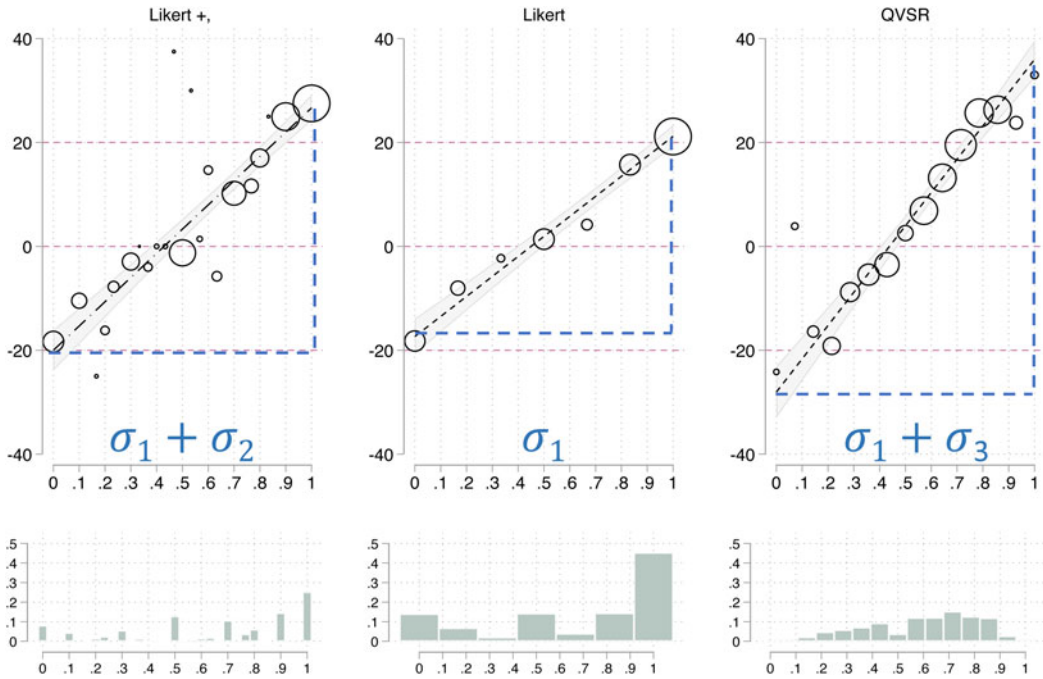


Figure 2. Donation to Gun-Related Advocacy Group and Responses to Gun Control Item. Binned scatterplot with a linear fit line. *Y-axis:* Average/predicted donation amount. *X-axis:* Survey answers by survey method, normalized to vary from 0 to 1. Survey item used: [gunC]. *Interpretation:* Scatterplot represents the average donation for respondents with the same X value, i.e., $E(Y|X=x)$. Dot size is proportional to the number of observations. Likert, in the center of the figure provides the benchmark. A visual comparison indicates that the coefficient for Likert+ is only marginally larger than that for Likert. Notice the difference in slope between Likert and QVSR. The full estimates are available in Figure 3. Compare also the bunching in Likert and the variation recovered under Likert+ and QVSR. Sparsely populated bins in Likert+ mean that not all 23 bins are visible to the naked eye.

has 23 response categories ranging from -15 to $+15$. While QVSR has 21 response categories in theory (from -10 to $+10$), in practice, few people put more than 7 votes on the same issue, resulting in 15 response categories (from -7 to 7).¹⁸ To insure comparability, we normalize \hat{u}_{ik}^L , \hat{u}_{ik}^{L+} and \hat{u}_{ik}^{QVSR} such that the lowest possible answer corresponds to zero ($-3/-15/-7$ for Likert, Likert+ and QVSR respectively) and the highest possible answer to 1 ($3/15/7$).

As shown in the bottom panel of Figure 2, when preferences are measured using a Likert item, the distribution of answers to the gun control item is uni-modal: answers bunch on one extreme of the scale (i.e., strong support for gun control). This pattern is much less pronounced with Likert+, implying that, while most respondents strongly support gun control, not everyone believes this issue is personally important to them. Partly by design, responses in QVSR exhibit no such bunching patterns.¹⁹

More response categories and less bunching imply more information (i.e., higher entropy) for QVSR and Likert+ on the one hand than for Likert on the other.²⁰ If Likert+ and QVSR’s higher entropy is more than just noise then, when comparing individuals with a higher score to individuals with a lower score, the former’s behavior should signal more intense preferences than the latter’s. Put differently, if a response category is a bin, people in a bin with a higher value should

¹⁸A few respondents chose to vote 8 times or more for the same issue. To avoid presenting results from bins that only include very small numbers of observations, we re-coded the 8 votes or more answers into a 7 votes answer.

¹⁹See Quarfoot *et al.* (2017) for within-individual evidence on this de-bunching process.

²⁰Shannon entropy scores capturing this difference are provided in Appendix F

be, on average, more likely to take action than people in a bin with a lower value. Quantitatively, this implies a positive and monotonic relationship between ordinal response categories on the one hand, and the mean of the outcome of interest—conditional on the response category—on the other. We examine this expectation by regressing each of the behavioral outcomes described in Table 1 over the corresponding normalized survey response variable (X) interacted with a categorical variable identifying the method used:

$$Y_i = \sigma_0 + \mu_1 D_{i,Likert+} + \mu_2 D_{i,QVSR} + \sigma_1 X_i + \sigma_2 X_i D_{i,Likert+} + \sigma_3 X_i D_{i,QVSR} + \sigma_4 J_4 + \dots + \sigma_j J_j + \epsilon_i \quad (1)$$

where J_4, \dots, J_j are dummy variables that indicate membership in a block used for block randomization (see Appendix C for more details). Regression coefficients σ_1 , $\sigma_1 + \sigma_2$ and $\sigma_1 + \sigma_3$ can be interpreted as the difference between $E(Y/X=1)$ and $E(Y/X=0)$ for Likert, Likert+ and QVSR respectively. The better tool is the one with not only more variation (or higher entropy) but also more informative variation in the form of a larger difference between the two quantities of interest, that is, the one with a larger regression coefficient. Monotonicity is also key: in the next section, we assess it visually by plotting the average value of Y_i for all respondents with the same value for X .²¹

3. Results

Figure 2 plots average donations to the gun control charities by response to the gun control question, further broken down by survey instrument. The lines capture the three regression coefficients mentioned in the previous section (see Figure 3 for the actual estimates). As shown on this figure, the regression slope is larger for QVSR than for Likert. This means that individuals who choose the end-of-scale response categories in Likert end up de-bunching under QVSR in ways that align with their behavior on the donation task. Specifically, people who donate less choose, on average, smaller values in QVSR than people who donate more. This is captured by the magnitude of the regression slope: individuals who do not donate are no longer pulling the regression slope down by “sharing” the extreme response categories with people who care enough to donate. Comparing the regression coefficients, we can also see that, in this case, the discrimination achieved with QVSR better aligns with preference intensity than that achieved with Likert+. For all three survey tools, the relationship between response category and average behavior is monotonic. Exceptions are due to sparsely populated bins.

As Figure 2 (center panel) shows, Likert does recover some information about preference intensity, as proxied by donation behavior. People who “strongly oppose” (1) or “strongly favor” (0) gun control donate more dollars to an organization that advocates for their preferred policy outcome than people who only “oppose” or “favor” gun control. The benefits of Likert+ and QVSR is that they distribute people across more response categories in ways that are informative of average donation behavior (that is, more variation/less bunching, a larger coefficient and monotonicity).

Figure 3 presents the same analysis for all tasks. Specifically, it plots regression coefficients obtained using equation 1 for all Y s and corresponding X s described in Table 1. Note a few important differences in how X s were computed. For the donation outcomes (first two columns in Figure 3), we use the normalized values of the response variables (X). When predicting the number of characters written, we use the normalized *absolute* values of the response variables. Indeed, our outcome variable does not capture *what* was written about the bill (i.e., for or against), only the overall effort spent writing about it. When predicting

²¹We also checked for a non-parametric relationship and find that, for all methods, the standard errors do not allow us to rule out a monotonic relationship.




Figure 3. Regression Coefficients for all Behavioral Outcomes.
Interpretation: A switch from the smallest response category (0) to the largest (1) is associated with a σ increase in Y . The increase is equal to σ times the standard deviation of Y . For the letter writing tasks (Minimum wage and Abortion), the predictor is the normalized absolute value of the response variable. For the punishment task, the predictor is the normalized difference between the gun control and the border wall response variables. For details on each, see text. *Note:* Sample sizes for the Gun and Immigration donation tasks (wave 1) are double the size of the samples sizes for the other tasks (wave 2). As a result, effect sizes are more precisely estimated for these two tasks. For details on each task, see text.

punishment in the dictator games, we use the normalized *difference* between responses on gun control and responses on the border wall (See Table 1). Higher positive values indicate that one favors gun control more intensely than one opposes the wall. Higher negative values indicated that one favors the wall more intensely than one opposes gun control. Given that the Independent recipient in the dictator game was opposed to gun control and opposed to the wall, we examine whether larger differences predict a higher likelihood of punishing the Independent.

The higher the regression coefficient in Figure 3, the better a given tool is at distinguishing between respondents with high and low preference intensity (as proxied by task-specific behavior). Again, Likert’s performance is noticeable: in line with the claim that Likert items capture a mix of preference orientation and preference intensity, people with end-of-scale answers behave differently from others (in all cases, the coefficient is positive and substantively large). Overall, the main issue with this measurement tool is whether, on hyper-partisan issues, such as gun control or abortion, there are enough people who do *not* choose end-of-scale answers to identify who truly cares and who doesn’t (see Appendix F for response histograms).

While Likert+ appears to carry more information on preference intensity than Likert, its discriminatory power (as captured by $\sigma_1 + \sigma_2$) is statistically indistinguishable from Likert’s on all 6 outcomes. Overall Likert+ relative performance is far less consistent than QVSR’s. For wave 1 outcomes (donation to an advocacy group task), QVSR outperforms Likert both substantively and statistically. Due to smaller sample sizes, results for wave 2 tasks exhibit larger standard errors. Still, a comparison of regression coefficients suggests that QVSR is more informative of preference intensity than Likert:

on all 4 outcomes, QVSR coefficients are at least twice the size of those found with Likert. In contrast, the coefficients for Likert+ represent, relative to Likert, a 50 percent increase at best.

Because of QVSR’s budget constraint, for individuals who use all their credits, votes on one issue is a linear combination of votes on other issues. As a result, the error terms across outcome-specific (or covariate-specific) equations are likely correlated. As a robustness check, we consequently re-run the analyses underpinning Figure 3 and estimate seemingly unrelated regressions models that account for this correlation (Zellner, 1962). Table 2 reports differences in coefficient size between methods. The results remain unchanged.

The bottom row of Table 2 reports the F-statistics under the null-hypothesis that, within a data collection wave, the sum of all between-method differences is equal to 0. This allows us to compare the performance of methods within a wave. For both waves, the F-statistic is 3 to 4 times larger when comparing QVSR and Likert then it is when comparing Likert+ to Likert, further indicating that, relative to Likert, the information gained with QVSR is substantively larger than that gained with Likert+. Still, when comparing QVSR and Likert+, fewer observations in wave 2 mean we cannot reject the null of no differences between QVSR and Likert+ at conventional levels.

If QVSR, or even Likert+, convey information on preference intensity that is not captured by Likert, then a test of a theory where preference intensity is a theoretically relevant concept could be affected by the measurement tool used. Next, we examine this conjecture focusing on a longstanding debate in political science on the relationship between policy preferences and material self-interest.

4. Where theory and measurement meet

A common starting point when studying preference formation is to expect people to support policies that positively affect their economic conditions and oppose policies that negatively affect them. According to public opinion scholars, this expectation finds limited empirical support. Instead, to explain preference formation, researchers have emphasized non-economic modes of reasoning such as value-based or partisan motivated reasoning (Sears and Funk, 1990; Margalit, 2013; Cavaille, 2023). Still, when it comes to preference intensity and the likelihood of behaving in opinion-congruent ways, material self-interest likely plays a key role. For example, while both men and women might support equal pay for equal work out of fairness concerns,

Table 2. Differences in coefficient size (seemingly unrelated models)

	QVSR versus Likert	Likert+ versus Likert	QVSR versus Likert+	QVSR versus Likert	Likert+ versus Likert	QVSR versus Likert+
Gun	0.59*** (0.14)	0.20* (0.10)	0.38*** (0.14)			
Immigration	0.37*** (0.13)	0.19* (0.10)	0.17 (0.14)			
Minimum Wage				0.03 (0.24)	0.08 (0.19)	0.05 (0.28)
Abortion				0.68*** (0.21)	0.38** (0.17)	0.36* (0.22)
DG Punish (1)				0.46 (0.34)	0.39 (0.33)	-0.06 (0.35)
DG Punish (2)				0.98*** (0.37)	0.38 (0.33)	0.49 (0.37)
N	2,409	2,512	2,417	976	1,017	995
F test	28.14	7.87	8.24	8.50	3.26	1.21
Prob > F	0.000	0.005	0.004	0.004	0.071	0.271

*p < 0.05, **p < 0.01, ***p < 0.001. We replicate Figure 3 analysis using seemingly unrelated models. This table reports the interaction between the preference variable and a dummy variable identifying the survey methods used. For example, for the gun donation outcome, the difference between the coefficient for Likert and that for QVSR is equal to 0.59. Bottom row: F-test for the null-hypothesis that the sum of the coefficients is equal to 0.

when it comes to taking action, women will be more likely to do so than men, meaning that women have stronger preferences on this issue than men. This point has been made repeatedly by John Krosnick when discussing the related concepts of attitude extremity and importance (Howe and Krosnick, 2017, 328).

Somewhat surprisingly, empirical analyses of preference formation rarely emphasize the distinction between preference orientation and preference intensity. Yet, this distinction has implications for measurement strategy. If material self-interest is hypothesized to affect preference orientation, then a binary variable measuring support for a given policy should, a priori, be enough to test this argument. If material self-interest is hypothesized to affect preference intensity, then QVSR might be a better measurement strategy.

Figure 4 examines the implication of overlooking the importance of measurement when examining the role of material self-interest. It plots the relationship between gender on the one hand and support for gender equality in the workplace on the other, measured using Likert, Likert+ and QVSR. Notice how, in Likert (middle panel), there is very little variation in survey answers: most people appear to *strongly* support workplace gender equality. The additional information gained by switching from Likert to Likert+ is informative of respondents' gender: women are more likely than men to be in the highest response category. In QVSR, the de-bunching is more consequential and, unlike Likert+, there is a clear linear and monotonic relationship between the number of votes in QVSR and the percentage of women as a share of individuals who cast the same number of votes.

As Figure 5 shows, the same pattern emerges when comparing parental leave preferences and a measure of one's proximity to childbirth. In Appendix G, we show similar results for affirmative

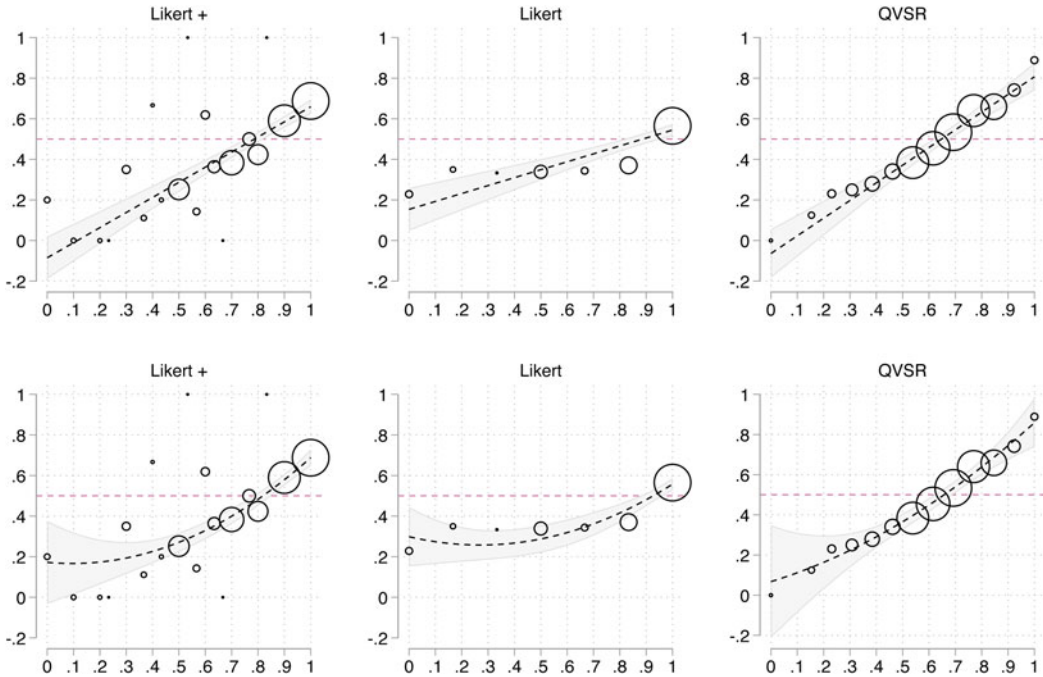


Figure 4. Respondent's Gender and Response to Pay Equity Item. Binned scatterplot with a linea (top panel) or quadratic (bottom panel) fit line. *Y-axis:* percent women in a given bin, predicted probability of respondent being a women (=1 if woman, 0 otherwise). *X-axis:* Survey answers by survey method, normalized to vary from 0 to 1. Survey item used: [sameS]. *Interpretation:* Scatterplot represents the share of women among respondents with the same X value, i.e., $E(Y|X=x)$. Dot size is proportional to the number of observations. Sparsely populated bins in Likert+ mean that not all 23 bins are visible to the naked eye.

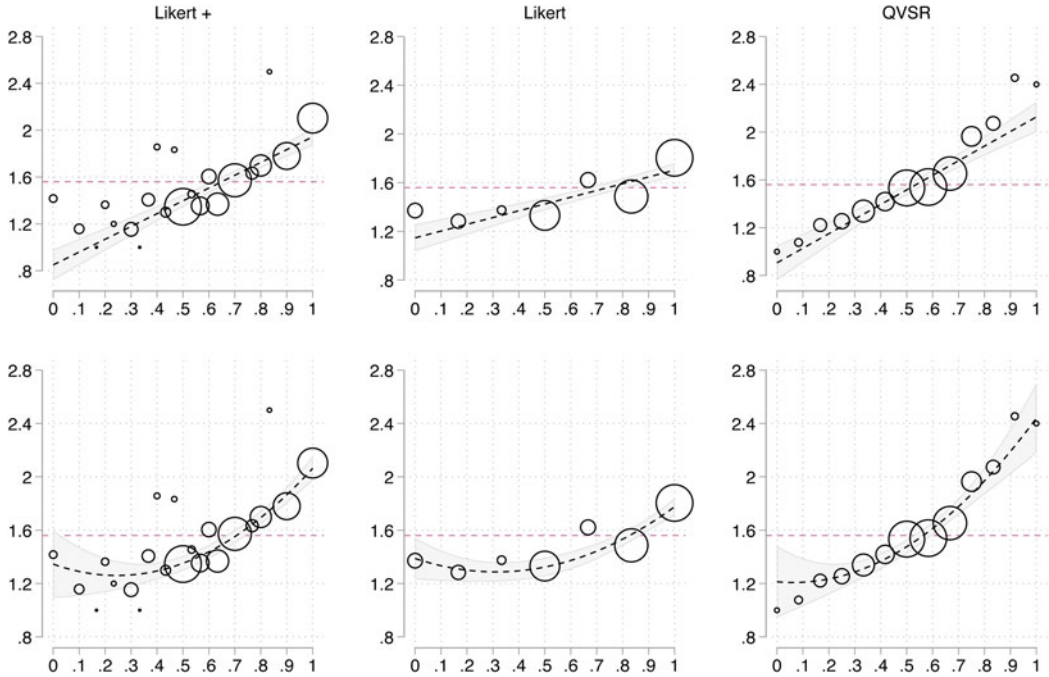


Figure 5. Respondent’s Proximity to Childbirth and Response to Parental Leave Item. Binned scatterplot with a linear or quadratic fit line. *Y-axis:* Average/predicted proximity to childbirth score (=1 if no young child and no plans to have any in future, =2 young children but no plans to have more, =3 if children planned or just had a child). *X-axis:* Survey answers by survey method, normalized to vary from 0 to 1. Survey item used: [paidL]. *Interpretation:* see note [Figure 4](#).

action and race, gun control and gun ownership, as well as increasing the minimum wage and the likelihood of benefiting from such increase. Because differences in preference intensity are imperfectly captured by Likert, using this item alone can produce the type of empirical patterns that have lead researchers to dismiss the theoretical relevance of material self-interest.

5. Conclusion

What do our argument and results imply for scholars interested in measuring preference intensity? Choosing a measurement strategy involves trade-offs between (1) maximizing interpretable variation, (2) minimizing survey costs, and (3) minimizing noise. QVSR performs well on (1) in the form of a larger number of better discriminating response categories. As we show, when it comes to documenting the importance of material self-interest, this can have substantive implications. QVSR does marginally worse on (2) in the form of longer survey time and more respondents dropping out. On (3), the improvement is minimal: standard errors remain similarly sized across methods, meaning that, as the quality of the signal increases, so does the noise, thus keeping the signal-to-noise ratio somewhat stable. In QVSR’s case, this could be due to the type of measurement error induced by a budget constraint that is too tight for some or too loose for others.

Additional work is thus needed to better understand where forced-choice methods like QVSR succeed and where they can be improved. For example, is the bulk of the work done by forcing respondents to consider issues jointly or does the quadratic pricing also play a key role? How might design tweaks —e.g., using a different cost function—help improve the noise-to-signal ratio? To answer these questions, future work might use linear rather than quadratic pricing

and compare QVSR to ranking methods.²² Relatedly, we have yet to examine the impact of changing the menu of options: would results differ had we included an item on the introduction of a wealth tax, or one on reparations? Within individuals, we would expect differences in the number of option-specific votes cast in one menu versus another. Still, how this will affect the cardinal information conveyed by the votes remains to be investigated. To facilitate such follow-up studies, we have made available a web application enabling researchers to vary QVSR's key features including pricing (e.g., linear versus quadratic) and the number of credits relative to the number of options.²³ We hope this will help spur future innovations in the measurement of policy opinions.

Ultimately, which measurement strategy to choose will depend on the type of financial constraints a researcher faces (e.g., survey time) as well as the type of policy issue being measured. When it comes to highly politicized issues, individual-level variance is much lower in Likert than in Likert+ and QVSR, which speaks in favor of QVSR. For less politicized issues, Likert+ might be enough. Possible menu effects are both a weakness and a strength of QVSR in particular, and forced-choice methods in general (e.g., conjoint analysis). On the one hand, they raise concerns about cross-study comparisons. On the other, they compel researchers to pick a menu of options that reflect theoretically relevant real-world constraints. This emphasis on theoretically-grounded design, while a weakness for exploratory research might be a strength in the deductive stage of a research project.

If there is one main take-away from our inquiry is that, faced with the expansion of survey-based research beyond descriptive public opinion polls, researchers need to take measurement seriously. Disciplinary boundaries have made it difficult: to the best of our knowledge, this is the first study to systematically compare and contrast measurement strategies derived from two distinct conceptualizations of human cognition and behavior, social psychology (in the case of Likert/Likert+) and economics (in the case of QVSR). We hope our conceptual and theoretical framework (see Appendix B for the more extensive discussion) will help future scholarship better specify their quantities of theoretical interest and identify the tools and strategies best adapted to measuring them.

Supplementary material. The supplementary material for this article can be found at <https://doi.org/10.1017/psrm.2024.27>. To obtain replication material for this article, <https://doi.org/10.7910/DVN/W27GR5>

Acknowledgements. We would like to thank Jonathan Ladd, Matthew Levendusky, Thomas Leeper, James Druckman, Samara Klar, Adam S. Levine, Gregory Huber, Kosuke Imai, Rich Nielsen and Michele Margolis for extensive help in the design phase of the pilot. Alisha Holland, Scott Page and Glen Weyl also provided important feedback. Presentations at the Institute for Advanced Study in Toulouse (IAST), Georgetown University, Columbia University, Oxford University, Princeton's CSDP and the University of Zurich generated many helpful comments. We would also like to acknowledge generous funding from the IAST inter-disciplinary seed grant and Carnegie's Bridging the Gap grant. Chen and van der Straeten acknowledge funding from the French National Research Agency (ANR) under the Investments for the Future (Investissements d'Avenir) program, grant ANR-17-EUR-0010.

References

- Abramson SF, Koçak K and Magazinnik A (2022) What do we learn about voter preferences from conjoint experiments?. *American Journal of Political Science* **66**, 1008–1020.
- Adida CL, Lo A and Platas MR (2018) Perspective taking can promote short-term inclusionary behavior toward Syrian refugees. *Proceedings of the National Academy of Sciences* **115**, 9521–9526.
- Bertrand M and Mullainathan S (2003) Enjoying the quiet life? Corporate governance and managerial preferences. *Journal of political Economy* **111**, 1043–1075.
- Bullock J and Lenz G (2019) Partisan bias in surveys. *Annual Review of Political Science* **22**, 325–342.

²²See footnote 8.

²³The resulting survey can be embedded into other online platforms, such as Qualtrics. This web application can be found at <https://qvsr.io>.

- Carsey TM and Layman GC** (2006) Changing sides or changing minds? Party identification and policy preferences in the American electorate. *American Journal of Political Science* **50**, 464–477.
- Cavaillé C** (2023) *Fair Enough? Support For Redistribution in the Age of Inequality*. Cambridge: Cambridge University Press.
- Cavaillé C, Chen DL and Van Der Straeten K** (2019) A decision-theoretic approach to understanding survey response: Likert versus quadratic voting for attitudinal research. *University of Chicago Law Review* **87**, 22–43.
- Ganter F** (2023) Identification of preferences in forced-choice conjoint experiments: reassessing the quantity of interest. *Political Analysis* **31**, 98–112.
- Hainmueller J, Hopkins DJ and Yamamoto T** (2014) Causal inference in conjoint analysis: understanding multidimensional choices via stated preference experiments. *Political analysis* **22**, 1–30.
- Hanretty C, Lauderdale B and Vivyan N** (2020) A choice-based measure of issue importance in the electorate. *American Journal of Political Science* **64**, 519–535.
- Hill SJ** (2022) *Frustrated Majorities: How Issue Intensity Enables Smaller Groups of Voters to Get What They Want*. Cambridge: Cambridge University Press.
- Howe L and Krosnick J** (2017) Attitude strength. *Annual Review of Psychology* **68**, 327–351.
- Krosnick J and Abelson R** (1992) The case for measuring attitude strength in surveys. In Tamur J (ed.), *Questions about Questions: Inquiries into the Cognitive Bases of Surveys*. New York: Russell Sage Foundation, pp. 177–203.
- Krosnick JA** (1991) Response strategies for coping with the cognitive demands of attitude measures in surveys. *Applied Cognitive Psychology* **5**, 213–236.
- Krosnick JA and Petty RE** (1995) Attitude strength: an overview. *Attitude Strength: Antecedents and Consequences* **1**, 1–24.
- Krosnick JA, Boninger DS, Chuang YC, Berent MK and Carnot CG** (1993) Attitude strength: one construct or many related constructs? *Journal of Personality and Social Psychology* **65**, 1132.
- Lalley SP and Weyl EG** (2018) Quadratic voting: How mechanism design can radicalize democracy. *AEA Papers and Proceedings* **108**, 33–37.
- Malhotra N, Krosnick JA and Thomas RK** (2009) Optimal design of branching questions to measure bipolar constructs. *Public Opinion Quarterly* **73**, 304–324.
- Margalit Y** (2013) Explaining social policy preferences: evidence from the great recession. *American Political Science Review* **107**, 80–103.
- Miller J and Peterson D** (2004) Theoretical and empirical implications of attitude strength. *The Journal of Politics* **66**, 847–867.
- Miller JM, Krosnick J and Fabrigar LA** (2017) The origins of policy issue salience. In Krosnick JA, Chiang I-CA and Stark TH (eds), *Political Psychology: New Explorations*. London: Routledge, pp. 125–171.
- Posner E and Weyl EG** (2018) *Radical Markets: Uprooting Capitalism and Democracy for a Just Society*. Princeton: Princeton University Press.
- Quarfoot D, von Kohorn D, Slavin K, Sutherland R, Goldstein D and Konar E** (2017) Quadratic voting in the wild: real people, real votes. *Public Choice* **172**, 283–303.
- Rabinowitz G and Macdonald SE** (1989) A directional theory of issue voting. *American Political Science Review* **83**, 93–121.
- Sauer C, Auspurg K, Hinz T and Liebig S** (2011) The application of factorial surveys in general population samples: the effects of respondent age and education on response times and response consistency. *Survey Research Methods* **5**, 89–102.
- Sears D and Funk C** (1990) The limited effect of economic self-interest on the political attitudes of the mass public. *Journal of Behavioral Economics* **19**, 247–271.
- Sides J, Tausanovitch C and Vavreck L** (2022) *The Bitter End: The 2020 Presidential Campaign and the Challenge to American Democracy*. Princeton, NJ: Princeton University Press.
- Wasow O** (2023) Text as behavior. http://omarwasow.com/text_as_behavior.pdf.
- Zaller J** (2012) What nature and origins leaves out. *Critical Review* **24**, 569–642.
- Zellner A** (1962) An efficient method of estimating seemingly unrelated regressions and tests for aggregation bias. *Journal of the American Statistical Association* **57**, 348–368.