A Decision Theoretic Approach to Understanding Survey Response: Likert vs. Quadratic Voting for Attitudinal Research

Charlotte Cavaillé†
Daniel L. Chen††
Karine Van der Straeten‡

"Likert items" are standard and widespread survey instruments. The most common version asks respondents to evaluate a statement by picking one of 5 (or 7) ordered responses ranging from strongly agree to strongly disagree. To interpret these answers, researchers often assume that respondents are sincere and report their true opinion with some random error. Yet, sincerity is not the only motive driving repondents' answers. Competing motives, most famously motives related to partisan identity, also matter. We propose a simple decision-theoretic model of survey answers that incorporates these different types of motives. We first show that respondents can systematically exagerate their views when asked about them using Likert items. We then show how, under certain conditions, Quadratic Voting for Survey Research (QVSR) can minimize this bias.

I. INTRODUCTION

At its heart, survey research tries to understand what individuals think and know about the world. Policymakers and social scientists conduct polls because they want to know what the public "thinks" on a given set of policy issues. This information can be used by policymakers to inform their decisions on specific policies. It is also used normatively by political scientists as a benchmark to judge the "quality" of democratic politics in a given policy arena.[1] For example, in the area of gun policy, political scientists rely on the gap between

---

    [1]   See generally, for example, Martin Gilens, *Affluence and Influence: Economic Inequality and Political Power in America* (Princeton 2012).

public opinion[2] and policy to conclude that interest groups, such as the National Rifle Association, must be "capturing" the policy process to impose their preferences on democratic majorities.[3] This line of work relies on the assumption that surveys successfully measure preferences on political or economic issues.

Most surveys rely on some variant of the Likert item, a survey technology developped in the 1930s by the psychologist Rensis Likert. Survey respondents are asked to evaluate a policy statement by picking one of five (or seven) ordered responses ranging from strongly agree/favor to strongly disagree/oppose. To interpret survey answers as a measure of "true" preferences, researchers need to assume that respondents are sincere when they pick a response category. Yet, there are good reasons to think that respondents' answers are shaped by additional motives and that these motives can conflict with respondents' desire to be sincere.

One motive examined by Professors John G. Bullock, Alan S. Gerber, Seth J. Hill, and Gregory A. Huber is partisan cheerleading.[4] They show that in surveys, strong partisans purposefully misreport facts about the state of the economy out of a desire to praise one party or criticize another. When this partisan motive is repressed (using monetary incentives to provide what one beliefs to be the true response), actors revert to answers that are closer to the "truth" and less predicted by their partisan identity.[5] These results indicate that when partisan cheerleading provides more utility than the sincerity motive, the observed answer will be different from respondents' true beliefs.

In a related paper, Professors Carlos Berdejó and Daniel L. Chen find that, in weeks prior to United States presidential

---

[2]   In the past three decades, an average of two-thirds of American citizens surveyed have indicated support for making gun control regulations stricter. See *Guns*, (Gallup 2018), archived at http://perma.cc/UK3V-EA6C.

[3]   In recent work on regulatory capture, Professor Daniel Carpenter similarly relies on public opinion polls to measure the influence of interest groups. He argues that political scientists first need to measure majority preferences and then look for evidence of interest groups moving policy away from these views to detect policy capture. This approach to survey data has been used in diverse areas, such as environmental, worker safety, food and drug, and labor regulations. Daniel Carpenter, *Detecting and Measuring Capture*, in Daniel Carpenter and David A. Moss, eds, *Preventing Regulatory Capture: Special Interest Influence and How to Limit It* (Cambridge 2014). For further analysis of the gap between public opinion polls and legislative action, see Gilens, *Affluence and Influence* (cited in note 1).

[4]   See generally John G. Bullock, et al, *Partisan Bias in Factual Beliefs about Politics*, 10 Q J Pol Sci 519 (2015).

[5]   Id at 563.

elections, judges on the US Courts of Appeals double the rate at which they dissent and vote along partisan lines.[6] It is unlikely that the judges' conception of a just decision has suddenly changed in the period preceding the election. What has changed is the possible returns of partisan cheerleading over sincerity. If highly experienced professionals making common-law precedent exhibit such a strong partisan bias, there is reason to believe that lay citizens answering political surveys (a comparatively low stake decision) should also, when faced with a conflict between signaling partisan and responding sincerily, deviate from expressing sincere beliefs or preferences.

In this article, we propose a simple and parsimonious decision-theoretic model of survey answers, which includes both the sincerity motive and other competing motives. We use this model to better understand how the match between true unobservable attitudes and observed survey answers varies with the survey instrument. We formally show that respondents are likely to systematically inflate their views when those are measured using Likert items.

We also examine an alternative survey technology, Quadratic Voting for Survey Research (QVSR). The QVSR instrument gives respondents a fixed budget to "buy" votes in favor or against a set of issues. Because the price for each vote is quadratic, it becomes increasingly costly to acquire additional votes to express support or opposition to the same issue. In a first exploration, Professor David Quarfoot and his co-authors compare QVSR to Likert-based survey instruments by randomly assigning respondents to one method or the other on M-Turk.[7] We complement this empirical approach by proposing a decision-theoretic analysis of the performance of QVSR versus Likert items.

The paper contributes to existing research on quadratic voting (QV). In its original formulation, QV is intended as a means of arriving at efficient social decisions when voting on policies that will be implemented.[8] Professor Steven P. Lalley and E. Glen Weyl primarily assume that influencing policy is

---

6   Carlos Berdejó and Daniel L. Chen, *Electoral Cycles Among U.S. Courts of Appeals Judges*, 60 J L & Econ 479, 486 (2017).

7   David Quarfoot, et al, *Quadratic Voting in the Wild: Real People, Real Votes*, 172 Pub Choice 283 (2016).

8   See E. Glen Weyl, *The Robustness of Quadrating Voting*, 172 Pub Choice 75, 79–81 (2017).

the main motivation of citizens.[9] In contrast, we model the implications of QV in a very different context: survey research. In addition, we explicitly model other potential motives such as a sincerity motive and a partisan motive.[10] We formally study how respondents solve the tradeoff between these potentially conflicting motives, depending on the survey instruments (Likert or QVSR). The model helps clarify the pros and cons of each survey intrument. Assuming the goal is to measure "true" attitudes, we lay out the conditions under which QVSR will outperforms Likert items.

II. A DECISION-THEORETIC MODEL OF SURVEY ANSWERS

Consider a number of policy issues, on which citizens may have any opinion between two extreme antagonistic positions. A survey is run to evaluate where the citizens stand on each of these issues. In this paper, extreme should not be understood in terms of how radical the proposition is. Instead, in the spirit of Likert items, extremity measure the extent to which one agrees with (supports) or disagrees with (opposes) a clearly defined proposal.

A.   Respondents' Motivation When Answering Surveys

We assume that an individual may have (at least) two (potentially) conflicting motives when answering the survey. On the one hand, she derives some intrinsic utility from reporting her "true opinion" on each issue. This might derive from some expressive benefits ("I am happy to tell what my opinion is"), or this might be induced by a psychological cost of lying. We call this motive the "sincerity motive." This is the motive generally assumed in the literature using survey data.[11]

---

9   Id.

10   Before moving on, one quick note to readers well versed in public opinion re-search: our emphasis is on the way in which survey technology affect the match between "true" and observed attitudes. In other words, unlike Professor John Zaller and scholars who build on his work, we are not examining the role motives such as partisan cheer-leading plays in shaping a respondent's underlying "true" attitude. Instead, we focus on the likelihood that this respondent will misreport it and how this likelihood might vary with the survey instrument. See generally, for example, Seth J. Hill, James Lo, Lynn Vavreck, and John Zaller, *How Quickly We Forget: The Duration of Persuasion Effects from Mass Communication*, 30 Pol Comm 521 (2013).

11   One assumption commonly made by social scientists using survey data is that they provide a faithful—if noisy—measure of respondents' views. See generally, for ex-ample, Christopher Achen, *Mass Political Attitudes and the Survey Response*, 69 Am Pol

On the other hand, we hypothesize that an individual may also care about how her answers will be read and interpreted by other people, which might conflict with this sincerity motive. This additional motivation might encompass a variety of psychological mechanisms, depending on the context and the question. For example, imagine that the government is considering whether a specific reform should be adopted or not, and that a survey is conducted to measure public support for or opposition to this reform. The respondent might be willing to strategically use her answers to the survey to influence policy making. Another motivation for the respondent might be to signal to herself, or to whoever is going to read the survey, that she has some socially desirable traits. For example, she may want to appear altruistic, non-racist, tolerant, etc. She might also want to signal a group identity. For example, if she is a Republican, and she expects Republicans to take specific positions on some issues, she may suffer a psychological cost from moving away from these typical "Republican positions." The partisan cheerleading mentioned by Professor Bullock and coauthors might thus be some combination of an identity motive (behave like a Republican) and an influence motive (get Republican policies passed). Whatever the source of this motivation, because of this "signaling motive," one position is particularly attractive to the respondent, which might be different from where she really stands.

B.   The Utility Function

To capture these two motivations (*sincerity* and *signaling*), we propose a simple general model describing how respondents answer surveys.[12] Assume a survey is run to measure respondents' views on K different policy issues. A position on any such issue is modeled as a real number in the interval $[-1, +1]$, where the extreme positions $+1$ and -1 denote perfect agreement with two opposite position on a given policy issue.[13]

Sci Rev 1218 (1975); Stephen Ansolabehere, Jonathan Rodden, and James M. Snyder Jr, *The Strength of Issues: Using Multiple Measures to Gauge Preference Stability, Ideological Constraint, and Issue Voting*, 102 Am Pol Sci Rev 215 (2008).

    [12]   Our model shares some similarities with that of John G. Bullock, et al, *Partisan Bias in Factual Beliefs about Politics*, 10 Q J Pol Sci 519 (2015) (cited in note 4), which studies systematic differences between Republican and Democrat voters in how they answer factual questions about economic facts.

    [13]   We here follow the common practice in survey research to ask respondents to place themselves on a single dimension. This dimension might run from full agreement

We assume that, on each issue $k = 1, \ldots, K$, respondent i is characterized by two parameters also lying in the interval $[-1, +1]$, her "true" opinion, denoted by x_{ik}, and the opinion she finds the most attractive because of the signaling motive, denoted by t_{ik}. We call the latter her "signaling target."

We assume that the utility a respondent derives from answering the survey, denoted by V_i, depends on her answers to the survey (her reported policy positions), denoted by $\hat{x}_i = (\hat{x}_{i1}, \ldots, \hat{x}_{iK}) \in [-1, +1]^K$, in the following way:

$$V_i(\hat{x}_i) = \sum_k [F_{ik}(\hat{x}_{ik}) + G_{ik}(\hat{x}_{ik})],$$

where functions F_{ik} and G_{ik} are single-peaked, and reach their maximum in $\hat{x}_{ik} = x_{ik}$ and $\hat{x}_{ik} = t_{ik}$ respectively. The first term in the utility function captures the *sincerity motive*. If for a given issue only this motive were present, the maximal utility an individual could get would be by reporting her true opinion on this issue. The second term on the right-hand side represents the *signaling motive*. If only this motive were present, the maximal utility an individual could get would be by reporting her signaling target. For the time being, we make no additional assumptions on functions F_{ik} and G_{ik} and on the signaling targets t_{ik}. In what follows, we will consider in more detail two examples of particular interest, one where the signaling motive is induced by a desire to influence policy, and the other where it is induced by a partisan identity.

C.   Survey Technology

The survey technology constrains survey respondents by specifying the set of answers that are admissible, that is, the set of answers the respondents can choose from. For example, under standard Likert items, a respondent can pick any answer on a pre-determined number of ordered responses (for example, "strongly oppose," "somewhat oppose," "neither oppose nor support," "somewhat support," "strongly support"). Under Quadratic Voting, there is a maximum number of points that the respondent can use to answer, and the marginal cost of

---

with a statement to full disagreement with the same statement. Another dimension often used is full agreement with a given statement (the government should be the only provider of healthcare services) and full agreement with its exact opposite (the government should have no role to play whatsoever in the prevision of healthcare services).

moving away from the neutral answer (here 0) increases linearly with the distance to this neutral answer.

D. Optimization Problem

Individuals are assumed to choose answers $(\hat{x}_i)$ that maximize the utility function V_i, subject to the constraints on answers imposed by the survey technology. Equipped with this very simple model, we can predict how respondents will answer the survey. In particular, our interest will be to discuss whether these reported views are a good measure of the "true opinion" (x_i). In the next section we describe answers under Likert items, and then turn to the QVSR technology in the following section.

III. PROPERTIES OF OPTIMAL ANSWERS UNDER LIKERT ITEMS

We consider first the case of Likert items. For simplicity, we will ignore the fact that there are in general only a discrete number of answers the respondent can choose from ("strongly favor", "somewhat favor", etc.); we will instead assume that she can pick any number in the $[-1, +1]$ interval. In that case, the set of admissible answers is simply $[-1, +1]^K$ and the individual solves the following optimization program:

$$\max_{\hat{x}_i \in [-1,+1]^K} V_i(\hat{x}_i) = \sum_k [F_{ik}(\hat{x}_{ik}) + G_{ik}(\hat{x}_{ik})].$$

Denote by $\hat{x}_i^L = (\hat{x}_{i1}^L, \ldots, \hat{x}_{iK}^L)$ the solution of this program.

A. Properties of the Optimal Responses

It is straightforward to check that the optimal answer on issue k ($\hat{x}_{ik}^L$) lies somewhere between x_{ik} and t_{ik} (that is, in the interval $[x_{ik}, t_{ik}]$ if $x_{ik} \leq t_{ik}$ and in the interval $[t_{ik}, x_{ik}]$ otherwise). Otherwise, the respondent could simultaneously improve on both objectives. Where exactly she will locate between these two positions depends on the shape of the functions F_{ik} and G_{ik}.

1. Concave sub-utility functions.

In particular, if the functions F_{ik} and G_{ik} are both concave with $F_{ik}^{'}(x_{ik}) = G_{ik}^{'}(t_{ik}) = 0$, there is a strictly interior solution. With Likert items, individuals answer by compromising between

their two motives. Answers incorporate information about both x_{ik} and t_{ik}.

2. Convex sub-utility functions.

By contrast, if the functions F_{ik} and G_{ik} are both convex, then the objective V_i is convex in $\hat{x}_{ik}$ and the individual either truthfully reports her true opinion x_{ik} on this issue, or she reports her "signaling target," t_{ik}.[14]

B. Systematic Misreporting and What It Means for Often Used Common Measurement Error Approaches

One assumption commonly made by social scientists using political survey data is that surveys using Likert items provide a faithful—if noisy—measure of respondents' "true" views.[15] Most of the emphasis is on measurement error due to "format and survey context, errors made by respondents, and so on."[16] One suggestion is to minimize measurement error by constructing a scale that combines similar Likert items.[17] This assumes that the covariance in the error term across K issues of the same "issue area" is zero.

In contrast, our model highlights the fact that respondents may instead deviate from their true views *in systematic ways* due to the signaling motives. As a result, misreporting (or "error") will be correlated across items: combining these items into a scale or index will not address the bias introduced by these signaling motives.

To illustrate this point, we document attitudinal patterns that are inconsistent with the dominant view of attitudes being measured only with random error. We present electoral cycles in such Likert indices developed by Professors Stephen Ansolobehere, Jonathan Rodden, and James M. Snyder Jr to measure economic and moral attitudes.[18] We use the General Social Survey's date of interview and cluster standard errors by

---

[14]   More specifically, an individual chooses $\hat{x}_{ik}^L = x_{ik}$ if $F_{ik}(x_{ik}) - F_{ik}(t_{ik}) > G_{ik}(t_{ik}) - G_{ik}(x_{ik})$ and $\hat{x}_{ik}^L = t_{ik}$ otherwise.

[15]   See, for example, Achen, 69 Am Pol Sci Rev at 1221 (cited in note 11); Ansolabehere, Rodden, and Snyder, 102 Am Pol Sci Rev at 216–17 (cited in note 11).

[16]   Ansolabehere, Rodden, and Snyder, 102 Am Pol Sci Rev at 217 (cited in note 11).

[17]   Id.

[18]   Stephen Ansolabehere, Jonathan Rodden, and James M. Snyder Jr, *Purple America*, 20 J Econ Perspectives 97 (2006).

year of interview. Each subsequent figure presents specifications with a full set of quarter-to-elect dummy indicators omitting quarter 16 (so November–January after an election is the comparison group), and also controls for seasonality (January–March, April–June, July–September, October–December). Figure 1 shows that Democrats are systematically more culturally conservative two quarters after a presidential election and two quarters after a midterm election (May–July). Figure 2 shows that Republicans are more culturally liberal one quarter after these elections. In the appendix, the corresponding patterns for economic attitudes are less pronounced, with both groups being more economically liberal two quarters after midterms. Next we analyze group cohesion as it varies over the electoral cycle. We calculate the average standard deviation in responses to each question for Democrats and for Republicans. Figure 4 shows that two quarters after elections, Republicans have more within-group cohesion on cultural issues. Figure 3 shows a similar pattern for Democrats two quarters after presidential elections. In the appendix, the pattern for economic attitudes are less systematic, though patterns still appear two quarters after elections (Fig. A1–A4). These patterns seem to contradict the hypothesis of random errrors.

When political opinions are measured using Likert items, there are a number of issues on which some respondents are likely to systematically misreport their true views. What can be said about the direction of this deviation? Will individuals appear in their answers to be more or less extreme than what they really are? In full generality, this deviation can go in any direction, depending on the relative position of the true opinion (x_{ik}) and of the partisan target (t_{ik}). Indeed, if $x_{ik} \leq t_{ik}$ then respondents will appaear more extreme. In the next two sections, we study two such situations where "systematic exaggeration" is likely to occur, and discuss wether quadratic voting might help alleviate this problem.

IV. THE "POLICY INFLUENCE MOTIVE"

One possible systematic exaggeration if when an individual wants to influence the decisions made by the government on issue k. The government is more likely to pay attention if surveys indicate full support for the reform. In other words, the target is $t_{ik} = +1$ if $x_{ik} > 0$, and $t_{ik} = -1$ if $x_{ik} < 0$, and there will be a strategic inflation in the reported intensity. In the polar

case where this policy influence motive is predominant, respondents will bunch at the extreme points of the Likert item.

A. Assumptions

To be more specific about the context, assume that K independent binary decisions have to be made by the government, say, implement a given reform or keep the status quo. In that case, $x_{ik} \in [-1, +1]$ is to be interpreted as the utility derived by individual i if reform k is implemented (compared to the status quo). Assume that a survey is run to evaluate the total utility that the implementation of each of the K reforms is likely to generate. We assume that the signaling part of the utility function has the following form: $G_{ik}(\hat{x}_{ik}) = x_{ik} S_{ik}(\hat{x}_{ik})$, where $S_{ik}(\hat{x}_{ik})$ is the probability that the reform is implemented if the individual reports $\hat{x}_{ik}$ (with $S'_{ik} > 0$). Note that in the strategic signaling motive, this "influence function" (S_{ik}) is weighted by how much the respondent is impacted by the reform (x_{ik}). To derive some simple closed-form solutions, we make the following assumptions:

$$F_{ik}(\hat{x}_{ik}) = -\frac{1}{2}\gamma_{ik}(x_{ik} - \hat{x}_{ik})^2 \text{ (quadratic sincerity motive)},$$

$$S_{ik}(\hat{x}_{ik}) = \sigma_{ik} \times \hat{x}_{ik} \text{ (linear policy influence)},$$

where parameter σ_{ik} captures the marginal impact of $\hat{x}_{ik}$ on the decision making process, and parameter $\gamma_{ik} \geq 0$ describes the weight of the sincerity versus signaling motive on issue k.

B. Optimal Responses under Likert

Under Likert, it is easy to check that the optimal answers are in that case:

$$\hat{x}_{ik}^L = \text{sign}(x_{ik}) \times \min\left[\left(1 + \frac{\sigma_{ik}}{\gamma_{ik}}\right)|x_{ik}|, 1\right], \tag{1}$$

where $sign(x_{ik}) = +1$ if $x_{ik} > 0$ and $sign(x_{ik}) = -1$ if $x_{ik} < 0$. Expression (1) shows that the optimal answer has the same sign as the 'true preference' (no misreporting in the direction of the preferences), but the intensity is always exaggerated. The size of the exaggeration is increasing with the ratio $\frac{\sigma_{ik}}{\gamma_{ik}}$. When this ratio

becomes large enough, the individual will choose to locate at one of the extremities of the [-1 , 1] scale. When such bunching occurs (in particular when only the policy influence motive is present ($\gamma_{ik} = 0$)), the only information that can be learned with the Likert technology is the direction of the preference; nothing can be learnt about intensity. Note that this is the situation originally motivating the use of quadratic voting in the seminal work of Lalley and Weyl.[19]

C. Optimal Responses under QVSR

One solution to this problem of strategic exaggeration might be to make reporting extreme values more costly than reporting moderate values. This is the basic idea underlying QVSR.[20] Formally, assume that the set of feasible answers under QVSR is:

$$\left\{\hat{x}_i = (\hat{x}_{i1}, \ldots, \hat{x}_{iK}) \in [-1, +1]^K : \sum_k \hat{x}_{ik}^2 \leq B\right\},$$

where $B \in \mathbb{R}_+$.

Deriving the optimal answers under QVSR is more complicated since it involves solving a constrained maximization program. The details of the proof are relegated in a technical appendix. As one can check in the appendix, the optimal response on issue k under QVSR is:

$$\hat{x}_{ik}^{QV} = \text{sign}(x_{ik}) \times \min\left[\frac{1}{1+\frac{2\lambda_i^*}{\gamma_{ik}}}\left(1 + \frac{\sigma_{ik}}{\gamma_{ik}}\right)|x_{ik}|, 1\right], \tag{2}$$

where λ_i^* is the Lagrange multiplier at the optimum. If $\sum_k \hat{x}_{ik}^2 \leq B$, meaning that optimal answers under Likert are within the QVSR budget set, then $\hat{x}_i^{QV} = \hat{x}_i^L$ and $\lambda_i = 0$. If $\sum_k \hat{x}_{ik}^2 > B$, optimal answers under Likert are not admissible under QVSR, and the individual has to report less extreme views.

---

[19]   See Weyl, 172 Pub Choice at 79–81 (cited in note 8).
[20]   Id.

D. Relative Performance of Likert and QVSR

As intuition suggests, the relative performance of Likert vs QVSR depends on the relative strength of the sincerity motive and policy influence motive. If the policy influence motive is very weak compared to the sincerity motive (i.e. $\frac{\sigma_{ik}}{\gamma_{ik}}$ close to 0), Likert items provide a good measure of preferences (see (1)). Indeed, reported views will be close to true opinions, with little bunching at extreme positions on the ordered responses. In that case, QVSR is not needed, and will even undermine the quality of the measure of preferences, since the binding budget constraint will prevent some respondents to report their true preferences.

By contrast, if the policy influence motive is strong enough, Likert will provide a poor measure of the intensity of preferences, because strategic considerations will induce respondents to bunch at extreme values. In that case, QVSR might represent a substantial improvement over Likert. Indeed, by making extreme reports more costly, it decreases the bunching at extreme positions observed with Likert, and is thus likely to generate better quality information about the intensity of preferences.

V. The "Partisan Consistency Motive"

A. Assumptions

When answering political surveys, the policy influence motive is not the only motive that may induce respondents to distort their true preferences. Another interesting example is a situation where citizens have strong partisan identities, and where even if they disagree with their preferred party's position on a specific issue, they suffer a psychological cost from reporting a divergent opinion. We will call this motive the "partisan consistency motive."[21] In that case, their signaling target t_{ik} on an issue is the position of the party with which they identify.

---

[21] When one's preferred party is not in power, the partisan consistency motive can additively interact with the policy influence motive: one might overreport agreement with abortion rights both because it is the right thing to do as a Democrat and because one wants to influence Republican policy on this issue.

Imagine a situation where party elites are very polarized, and consider an individual who generally agrees with her preferred party regarding the "direction" of the policy (that is, x_{ik} and t_{ik} have the same sign on most issues), but who is generally less extreme ($|x_{ik}| \leq |t_{ik}|$ on most issues). Under Likert, such an individual, because she wants to look like a 'good Republican' or like a 'good Democrat', will pick more extreme answers than she would if she were just reporting truthfully her own opinion.

To derive some simple closed-form solutions, we study a simple example with quadratic sub-utility functions for both the sincerity and the signaling motives. We assume that the utility function V_i is:

$$V_i(\hat{x}_i) = -\frac{1}{2} \sum_k \alpha_{ik} [(1 - \beta_{ik})(\hat{x}_{ik} - x_{ik})^2 + \beta_{ik}(\hat{x}_{ik} - t_{ik})^2],$$

with $\alpha_{ik} > 0$ and $\beta_{ik} \in [0,1]$. Parameter α_{ik} is the importance put on issue k when answering the survey, and parameter β_{ik} is the relative weight of the partisan consistency motive compared to the sincerity motive for issue k.

B.   Optimal Responses under Likert

Under Likert technology, it is easy to check that the solution of this optimization program is:

$$\hat{x}_{ik}^L = (1 - \beta_{ik})x_{ik} + \beta_{ik}t_{ik}. \tag{3}$$

If $\beta_{ik} = 0$ (only the sincerity motive is active), the individual has no incentive to misreport her view, and $\hat{x}_{ik}^L = x_{ik}$. But as soon as $\beta_{ik} > 0$, the individual has the incentive to move away from her true opinion in the direction of her partisan target.

Note that under Likert, how much the individual values her answer to this question compared to other questions in the survey (parameter α_{ik}) does not influence her answers. Indeed, each question is treated in isolation.

C. Optimal Responses under QVSR

Under QVSR technology, we show in the appendix that the optimal response on issue k under QVSR is:

$$\hat{x}_{ik}^{QV} = \frac{1}{1+2\frac{\lambda_i^*}{\alpha_{ik}}}[(1-\beta_{ik})x_{ik} + \beta_{ik}t_{ik}], \tag{4}$$

where λ_i^* is the Lagrange multiplier at the optimum.

D. Relative Performance of Likert and QVSR

Note first that, as in the case of the policy influence motive, if the partisan motive is very weak compared to the sincerity motive (i.e. β_{ik} close to 0), Likert items provide a good measure of preferences (see (3)). In that case again, QVSR is not needed, and will even undermine the quality of the measure of preferences, since the binding budget constraint will prevent some respondents to report their true preferences.

Consider now cases where the partisan motive can be potentially strong. As soon as the budget constraint is binding, compared to Likert, QVSR "shrinks" all answers towards the neutral answer (0). Expression (4) shows that this "contraction" can be heterogenous across issues: more points will be given to issues with a higher α_{ik}, meaning that respondents will put more points on issues that they consider as important. In that case, the relative performance of QVSR vs. Likert at measuring "true" opinions depends on the statistical relationship between α_{ik} (the importance of the issue) and β_{ik} (the relative importance of the partisan motive compared to the sincerity motive). Depending on this relationship, either method can dominate.

If, for an individual, the α_{ik} are the same for all issues, expression (4) shows that the optimal answers under QVSR are given, compared to Likert, by just 'shrinking' all answers proportionally towards the neutral answer (0), until one satisfies the budget constraint.[22] If the partisan targets are more extreme than the respondents' true views ($|t_{ik}| > |x_{ik}|$), QVSR will move answers in the correct direction (compared to Likert). Yet, it is important to note that QVSR will not "purge" reported answers of the partisan motive: answers will still be a convex combination of the true opinion and the partisan target, with

---

[22] More specifically, using condition (6) in the appendix, one can check that: $\hat{x}_{ik}^{QV} = \sqrt{\frac{B}{\Sigma_m(\hat{x}_{im}^L)^2}} * \hat{x}_{ik}^L$.

exactly the same relative weights as under Likert. In that sense, QVSR will not perform better than Likert.

If high α_{ik} tend to be associated with low β_{ik}, more votes will be put on issues with a strong sincerity motive, and QVSR might perform better than Likert at measuring "true opinions." There are reasons to expect such a positive correlation between the importance of the issue and the strength of the sincerity motive. Indeed, consider an individual who cares strongly about some issues, and considers others as secondary or not very important. On the former set of issues, the individual will be ready to collect information, invest some time and effort to think about the pros and cons of various policies, and eventually form a strong, independent opinion. For such issues, the sincerity motive is likely to be strong and the partisan motive weak. By contrast, consider the issues in the latter set. Such issues are issues the individual does not really care about and has not thoughtfully reflected upon. In that case, she might be happy to use the party line as the main determinant of her answers. To make this argument more clearly, consider the extreme case where there are two types of issues: those about which the individual cares and where the sincerity motive is predominant, say $\alpha_{ik} = 1$ and $\beta_{ik} = \varepsilon$, with $\varepsilon << 1$, and those about which the individual does not care and where the partisan motive is predominant, say $\alpha_{ik} = \varepsilon$ and $\beta_{ik} = 1 - \varepsilon$. In that case, under Likert, the individual will report her true opinion on the first set of issues, and will report her partisan target on the second set (See (3)). Under QVSR, assuming that the budget constraint is sufficiently binding, she will put no points on the second set of issues, and she will allocate all her points on the issues with a strong sincerity motive. In such a situation, QVSR is likely to represent a significant improvement over Likert.

Note last that if high α_{ik} tend to be associated with high β_{ik}, the exact opposite argument will prevail and QVSR might perform worse than Likert.

VI. CONCLUSION

We have good reasons to think that systematic biases in translating true into observed attitudes exist. We formally study how conflicting motives (more specifically, a sincerity motive, a policy-influence motive, and a partisan motive) may induce some tradeoffs when individuals answer survey questions. We show that QVSR performs better than Likert items on average when

this bias generates bunching toward one or both ends of the ordered response categories, or when across issues, there is a positive correlation between the strength of the sincerity motive and the general importance of the question to the respondent. Yet, we also highlight several limits, which indicate that a specific set of assumptions need to be tested before being able to decide, for each specific context, which instrument dominates.

VII. TECHNICAL APPENDIX

A.1 Optimal Responses under QVSR: The "Policy Influence" Case

Consider the "policy influence motive" case, with the following utility function:

$$V_i(\hat{x}_i) = \sum_k \left[-\frac{1}{2} \gamma_{ik} (\hat{x}_{ik} - x_{ik})^2 + x_{ik} \sigma_{ik} \hat{x}_{ik} \right].$$

Under QVSR, the individual maximizes her utility V_i subject to the budget constraint $\sum_k \hat{x}_{ik}^2 \leq B$. Write the Lagrangian as:

$$\mathcal{L}(\hat{x}_i, \lambda_i) = \sum_k \left[-\frac{1}{2} \gamma_{ik} (\hat{x}_{ik} - x_{ik})^2 + x_{ik} \sigma_{ik} \hat{x}_{ik} \right] + \lambda_i \left(B - \sum_k (\hat{x}_{ik})^2 \right),$$

where λ_i is the Lagrange multiplier. Taking the derivatives with respect to $\hat{x}_{ik}$, one can check that the first order condition gives:

$$\hat{x}_{ik} = \begin{cases} +1 \text{ if } \left(1 + \dfrac{\sigma_{ik}}{\gamma_{ik}}\right) \times |x_{ik}| > 1 + \dfrac{2\lambda_i}{\gamma_{ik}} \text{ and } x_{ik} > 0, \\[2ex] -1 \text{ if } \left(1 + \dfrac{\sigma_{ik}}{\gamma_{ik}}\right) \times |x_{ik}| > 1 + \dfrac{2\lambda_i}{\gamma_{ik}} \text{ and } x_{ik} < 0, \\[2ex] \dfrac{1 + \dfrac{\sigma_{ik}}{\gamma_{ik}}}{1 + \dfrac{2\lambda_i}{\gamma_{ik}}} \times x_{ik} \text{ otherwise.} \end{cases}$$

If $\sum_k \left(\hat{x}_{ik}^L\right)^2 \leq B$: responses are the same as under Likert (the budget constraint is not binding, and at the optimum the Lagrange multiplier is equal to 0).

If $\sum_k \left(\hat{x}_{ik}^L\right)^2 > B$, then satisfying the budget constraint implies that:

$$\sum_{k:x_{ik}>0} \left(\min\left(\frac{1+\frac{\sigma_{ik}}{\gamma_{ik}}}{1+\frac{2\lambda_i}{\gamma_{ik}}} \times x_{ik}, +1\right)\right)^2 + \sum_{k:x_{ik}<0} \left(\max\left(\frac{1+\frac{\sigma_{ik}}{\gamma_{ik}}}{1+\frac{2\lambda_i}{\gamma_{ik}}} \times x_{ik}, -1\right)\right)^2$$
$$= B. \qquad (5)$$

Note that the left-hand side of the equality is strictly decreasing in λ_i. It takes the value $\sum_k \left(\hat{x}_{ik}^L\right)^2 > B$ when $\lambda_i = 0$, and it converges towards 0 as λ_i goes to $+\infty$. Therefore, there exists a unique positive λ_i such that equality (5) is satisfied. Note that this value depends on all the parameters $\gamma_i = (\gamma_{i1}, \ldots, \gamma_{iK})$, $\sigma_i = (\sigma_{i1}, \ldots, \sigma_{iK})$ and on the true preferences $x_i = (x_{i1}, \ldots, x_{iK})$. Denoting the Lagrange multiplier at the optimum by λ_i^*, under QVSR, the optimal response on issue k is therefore:

$$\hat{x}_{ik}^{QV} = \text{sign}(x_{ik}) \times \min\left[\frac{1+\frac{\sigma_{ik}}{\gamma_{ik}}}{1+\frac{2\lambda_i^*}{\gamma_{ik}}}|x_{ik}|, 1\right],$$

which is expression (2) in the main text.

A.2 Optimal Responses under QVSR: The "Partisan Consistency" Case

Consider the "partisan consistency motive" case, with the following utility function:

$$V_i(\hat{x}_i) = -\frac{1}{2}\sum_k \alpha_{ik}[(1-\beta_{ik})(\hat{x}_{ik} - x_{ik})^2 + \beta_{ik}(\hat{x}_{ik} - t_{ik})^2].$$

Under QVSR technology, the individual now solves the following optimization program:

$$\max_{\hat{x}_i \in [-1,+1]^K} -\frac{1}{2}\sum_k \alpha_{ik}[(1-\beta_{ik})(\hat{x}_{ik} - x_{ik})^2 + \beta_{ik}(\hat{x}_{ik} - t_{ik})^2] + \lambda_i(B - \sum_k (\hat{x}_{ik})^2),$$

where λ_i is the Lagrange multiplier. First order condition with respect to $\hat{x}_{ik}$ now yields:

$$\hat{x}_{ik}^{QV} = \frac{\alpha_{ik}}{\alpha_{ik} + 2\lambda_i}[(1-\beta_{ik})x_{ik} + \beta_{ik}t_{ik}] = \frac{\alpha_{ik}}{\alpha_{ik} + 2\lambda_i}\hat{x}_{ik}^L.$$

If $\sum_k \left(\hat{x}_{ik}^L\right)^2 \leq B$: responses are the same as under Likert (the budget constraint is not binding).

If $\sum_k \left(\hat{x}_{ik}^L\right)^2 > B$, then satisfying the budget constraint implies that:

$$\sum_k \left(\frac{\alpha_{ik}}{\alpha_{ik}+2\lambda_i} \hat{x}_{ik}^L\right)^2 = B. \tag{6}$$

Note that the left-hand side of the equality is strictly decreasing in λ_i, taking the value $\sum_k \left(\hat{x}_{ik}^L\right)^2$ strictly higher than B when $\lambda_i = 0$, and is converging towards 0 as λ_i goes to $+\infty$. Therefore, there exists a unique positive λ_i such that equality (6) is satisfied. Denoting the Lagrange multiplier at the optimum by λ_i^*, the optimal response on issue k under QVSR is therefore:

$$\hat{x}_{ik}^{QV} = \frac{1}{1 + 2\dfrac{\lambda_i^*}{\alpha_{ik}}} \hat{x}_{ik}^L,$$

which is expression (4) in the main text.

FIGURE 1: MORAL ATTITUDES, DEMOCRATS
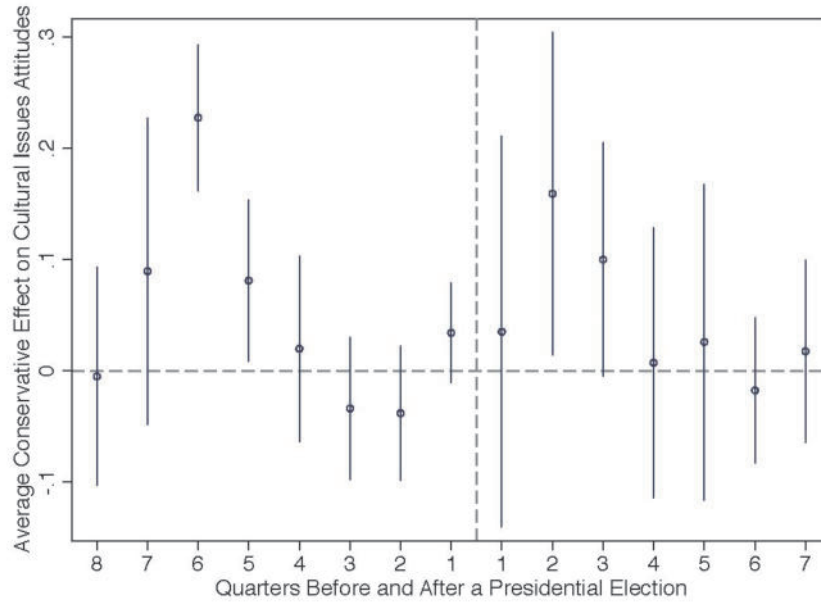
FIGURE 2: MORAL ATTITUDES, REPUBLICANS
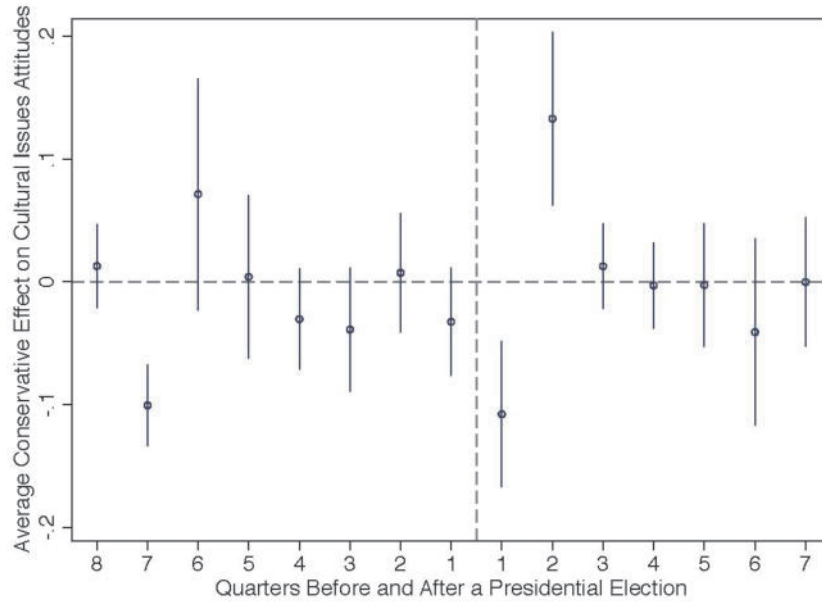
FIGURE 3: MORAL ATTITUDES (STANDARD DEVIATION), DEMOCRATS
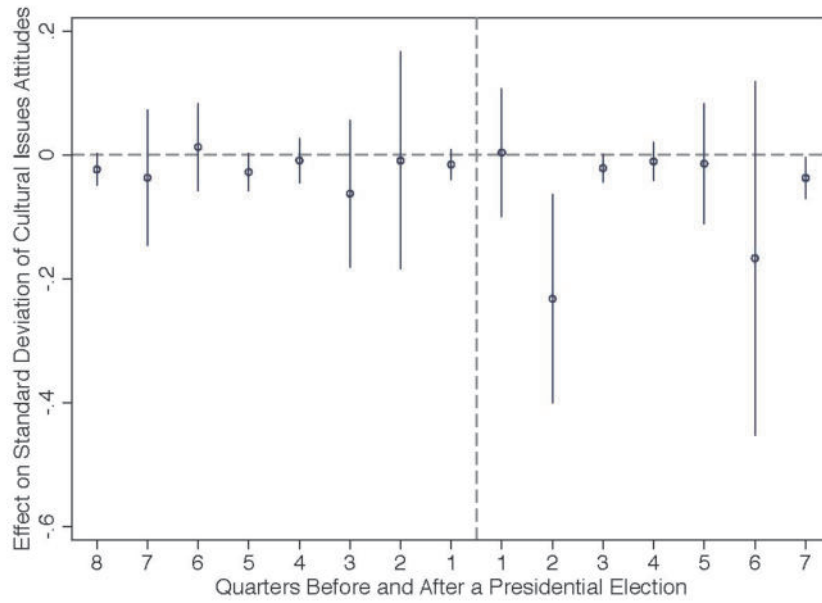
FIGURE 4: MORAL ATTITUDES (STANDARD DEVIATION),
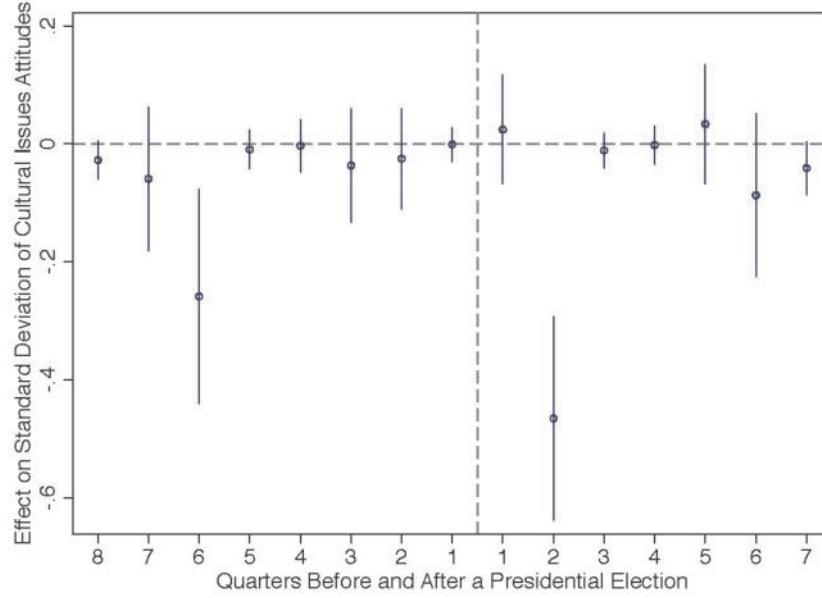REPUBLICANS

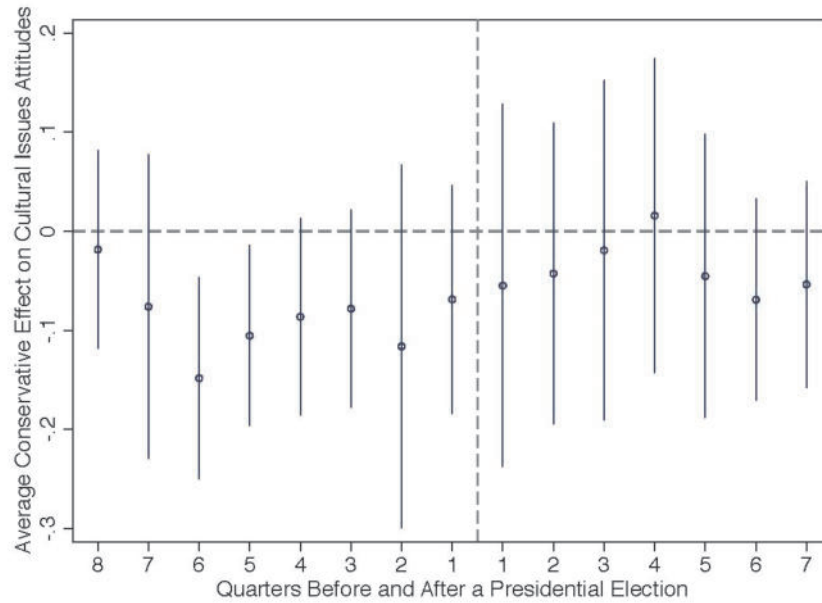

FIGURE A1: ECONOMIC ATTITUDES, DEMOCRATS

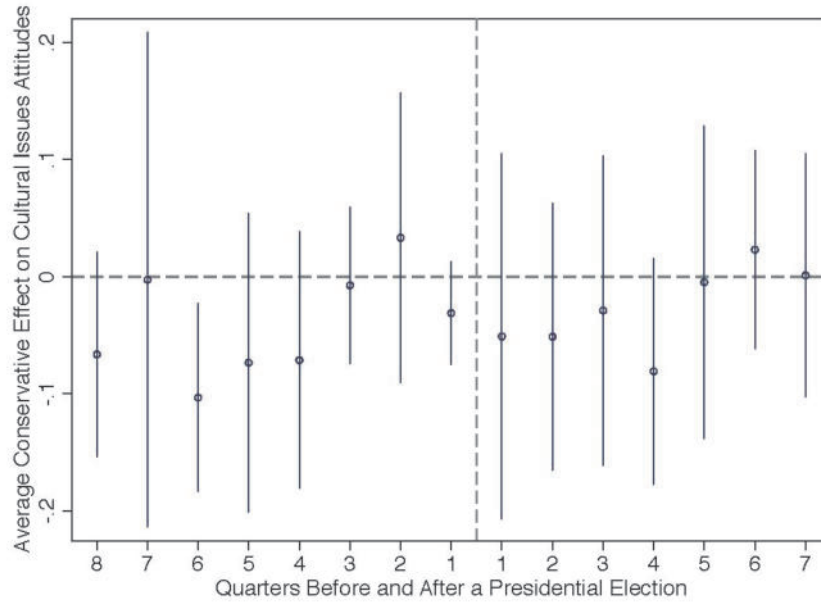FIGURE A2: ECONOMIC ATTITUDES, REPUBLICANS

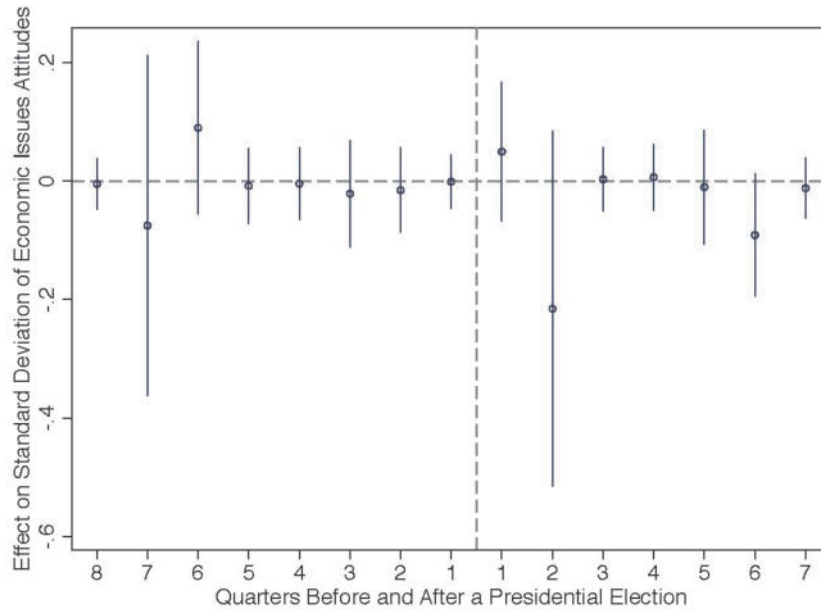FIGURE A3: ECONOMIC ATTITUDES (STANDARD DEVIATION), DEMOCRATS

FIGURE A4: ECONOMIC ATTITUDES (STANDARD DEVIATION),
REPUBLICANS