Advance Access publication September 29, 2015

Political Analysis (2015) 23:518–533
doi:10.1093/pan/mpv022

Closeness Counts: Increasing Precision and Reducing Errors in Mass Election Predictions

Kai Quek
Department of Politics and Public Administration, University of Hong Kong
e-mail: quek@hku.hk

Michael W. Sances
Department of Political Science, University of Memphis
e-mail: msances@memphis.edu (corresponding author)

Edited by R. Michael Alvarez

Mass election predictions are increasingly used by election forecasters and public opinion scholars. While they are potentially powerful tools for answering a variety of social science questions, existing measures are limited in that they ask about victors rather than voteshares. We show that asking survey respondents to predict voteshares is a viable and superior alternative to asking them to predict winners. After showing respondents can make sensible quantitative predictions, we demonstrate how traditional qualitative forecasts lead to mistaken inferences. In particular, qualitative predictions vastly overstate the degree of partisan bias in election forecasts, and lead to wrong conclusions regarding how political knowledge exacerbates this bias. We also show how election predictions can aid in the use of elections as natural experiments, using the effect of the 2012 election on partisan economic perceptions as an example. Our results have implications for multiple constituencies, from methodologists and pollsters to political scientists and interdisciplinary scholars of collective intelligence.

1 Introduction

Can ordinary citizens accurately predict election outcomes? This question is of growing importance in political science. From the standpoint of election forecasting, citizen forecasts of elections outperform predictions based on intended vote choice (Rothschild and Wolfers 2013). For public opinion scholars, citizen estimates of electoral viability provide insight into the power of partisan bias (Thibodeau et al. 2015; Daniller, Silver, and Moehler 2013) and the potential for collective wisdom (Page and Shapiro 1992; Murr 2009; Miller et al. 2012). Whether actors can anticipate electoral results also has implications for the use of elections as exogenous treatments that impact political and economic outcomes (Snowberg, Wolfers, and Zitzewitz 2007a; Gerber and Huber 2010; Caughey and Sekhon 2011).

Yet, while mass electoral predictions are important to numerous theoretical and empirical questions in political science, existing measures of predictions leave much to be desired. The reason is that such measures overwhelmingly ask about *who will win*, whereas researchers are typically interested in *expected voteshare*. For instance, in the field of election forecasting, scholars almost always express their forecasts in terms of expected voteshares. Yet, because respondent expectations are measured using a binary variable, analysts must impose additional functional form assumptions to obtain the desired continuous voteshare predictions.

Authors' note: We thank Eitan Hersh, Krista Loose, Michele Margolis, and Jim Snyder for comments. We also thank the MIT Political Experiments Research Lab for support. Sances thanks the Center for the Study of Democratic Institutions at Vanderbilt University for institutional support. This research received IRB approval from the Committee on the Use of Humans as Experimental Subjects at the Massachusetts Institute of Technology (protocol #1203004950). Replication data for this article may be found in Sances and Quek (2015). Supplementary Materials for this article are available on the Political Analysis Web site.

© The Author 2015. Published by Oxford University Press on behalf of the Society for Political Methodology.
All rights reserved. For Permissions, please email: journals.permissions@oup.com

https://doi.org/10.1093/pan/mpv022 Published online by Cambridge University Press

This measurement problem also affects studies of partisan bias in expectations, which require knowledge of the "correct" expectation. Yet having only two possible choices—either the Democrat or the Republican will win—precludes any objective measure of bias. For example, if 90% of Republicans believed Mitt Romney would win in 2012, it might also be the case that 90% of Republicans believed he would win by only 1 percentage point. The "who will win" question suggests massive partisan bias, while asking if voteshare would reveal this bias to be quite modest. Further, the reduction in variation that comes with a binary prediction measure limits our ability to analyze the individual-level covariates of respondent forecasts, such as political information.

Election outcomes are also often used as natural experiments. To be a valid natural experiment, voters and elites must be uncertain about the actual future voteshare. One way to validate this assumption is to show that voters and elites are imperfect predictors of election outcomes. Researchers might also seek to validate their results by testing whether the effects of elections are greatest among those most surprised by the outcome. Yet, once again, asking these actors who they think will win fails to pick up the level of surprise. If, as in the example above, 90% of Republicans believed Mitt Romney would win by a very slim margin, the true level of surprise is quite low; yet if these 90% believed Romney would win by a landslide, surprise is high, but undetectable without asking about voteshare itself.

Given the numerous drawbacks of existing measures, why do researchers not simply ask respondents to predict voteshare? At least part of the reason lies in survey researchers' aversion to asking respondents about quantities. Ansolabehere, Meredith, and Snowberg (2013) highlight surveys' overwhelming reliance on qualitative measures of quantitative information, such as economic perceptions and the likelihood of different events. These authors attribute the practice to researchers' belief that quantitative measures are too demanding for survey respondents, who may lack the ability or motivation to reason about numbers. Ansolabehere, Meredith, and Snowberg (2013) also show that these low expectations are in fact unwarranted: respondents can estimate gas prices and unemployment rates in their own states with some precision.

In this article, we test whether ordinary citizens can make accurate quantitative predictions of future election outcomes, and show how conventional qualitative measures can lead to severely mistaken inferences. We first demonstrate that respondents can indeed produce sensible estimates of voteshare. On average, respondents in our survey predicted the result of the 2012 presidential election within 1.1 percentage points of the actual outcome. This level of accuracy rivals forecasts of the 2012 election made by professional forecasters.

Having demonstrated that respondents can handle numerical predictions, we then show how traditional qualitative forecasts give misleading results. We begin with an examination of partisan bias, showing that qualitative predictions vastly overstate the degree of partisan bias in election predictions. Moreover, qualitative predictions lead one to believe that political information exacerbates the bias, whereas quantitative predictions show information narrowing the prediction gap between partisans.

Last, we demonstrate the utility of quantitative predictions as an aid to election-based natural experiments. Using the 2012 presidential election as an example, we test for the effect of partisan control on economic perceptions. As a way of validating both our prediction measure and the natural experiment itself, we show that the effect of the election increases the further respondents' predictions strayed from the true voteshare. In contrast, this pattern is much more difficult to detect using the standard qualitative prediction measure.

2 The Promise of Mass Election Predictions

Predicting election outcomes has for many years been an active area of research in political science (Lewis-Beck and Rice 1992). While traditionally these forecasts have been made by experts, who apply sophisticated statistical techniques to polling and economic data, interest has been increasing in the use of forecasts by non-experts (Rothschild 2009). Many of these forecasts are made using election markets, where participants buy and sell future contracts based on their expectations of election outcomes. Reviewing four recent presidential elections, Wolfers and Zitzewitz (2004) show

https://doi.org/10.1093/pan/mpv022 Published online by Cambridge University Press

that market predictions in the final week of the campaign have an absolute error of 1.5 percentage points, whereas final Gallup polls are off by 2.1 percentage points (see also Erikson and Wlezien 2008, 2012).

While traders in prediction markets are likely non-representative of the general public, more recent research has shown that the mass public is also surprisingly competent at forecasting elections. Since 1952, the American National Election Studies (ANES) have asked respondents, "Who do you think will be elected President in November?" Rothschild and Wolfers (2013) show that for national election outcomes, these binary expectation measures rival the accuracy of predictions based on the ANES vote intention question. Moreover, mass expectations of state outcomes significantly outperform forecasts based on state vote intent proportions. Others find these expectation questions sometimes outperform prediction markets, both using respondents from the ANES (Graefe 2014) as well as an online convenience sample (Miller et al. 2012).

In addition to being a powerful and relatively inexpensive forecasting tool, expectations about election outcomes can also provide insight into the influence of partisanship on information processing (Campbell et al. 1960; Bartels 2002). In a study of motivated reasoning among partisans, Thibodeau et al. (2015) argue that partisans seek evidence suggesting their preferred candidate will win, and ignore evidence that their candidate will lose. To test this claim, they experimentally assign favorable or unfavorable election predictions, and find that these treatments only affect expectations when they favor the respondent's preferred candidate. Similarly, Daniller, Silver, and Moehler (2013) use observational data to show that partisan media consumption by citizens is associated with greater bias in election predictions.

Citizen predictions are also implicated in the use of election outcomes as natural experiments. For instance, regression discontinuity designs rely on the assumption that close election outcomes are as-if randomly assigned, which necessarily implies that voters (among other actors) cannot perfectly predict voteshare (Caughey and Sekhon 2011). Information about the predicted voteshares of political actors can therefore serve as a useful validity check in these designs (Eggers et al. 2015, 268–9), especially when voter behavior is itself the outcome of interest (Matsubayashi 2013).

In addition to discontinuity designs, researchers have long been interested in exploiting the panel nature of elections to test for effects on political and economic behavior. At the individual level, Gerber and Huber (2010) use the 2006 midterm election outcome to test whether partisan economic perceptions—the tendency for members of the governing party to view the economy more favorably than members of the opposing party—are causally related to political events. Others have used panel surveys to test for an effect of election outcomes on trust in government (Ginsberg and Weissberg 1978; Keele 2005), political efficacy (Finkel 1985), and perceptions of electoral fairness (Moehler 2009; Sances and Stewart 2015). At a more macro level, several studies in political economy have attempted to estimate the effects of election outcomes on stock prices (Herron 2000; Knight 2006; Snowberg, Wolfers, and Zitzewitz 2007a, 2007b). As with regression discontinuity designs, the validity of these panel designs is bolstered when researchers can show that election outcomes are indeed a surprise for some respondents, as well as that treatment effects are stronger where outcomes are more surprising (Snowberg, Wolfers, and Zitzewitz 2007a; Margolis 2015).

3 Limitations of Existing Prediction Measures

Citizen predictions of election outcomes have strong potential for a variety of applications. However, existing measures of citizen predictions fail to fully realize this potential. The reason is that existing measures ask respondents who they think will win, rather than the winner's vote margin. For instance, the ANES asks, "Who do you think will be elected President in November?" and gives respondents the option of choosing one of two options. This practice not only severely reduces the amount of information available to researchers, but also reflects a disconnection between theory and measurement. In the forecasting literature, for example, researchers are ultimately interested in the winner's vote margin. To make the leap from a binary measure of who will win to a continuous measure of predicted voteshare, researchers must make several additional assumptions (Lewis-Beck and Tien 1999).

https://doi.org/10.1093/pan/mpv022 Published online by Cambridge University Press

For instance, Rothschild and Wolfers (2013) use a probit specification to translate the binary prediction question from the ANES into continuous predicted voteshare. They model the probability of predicting a Democratic win as a function of a latent variable, the respondent's expectation of Democratic voteshare. This latent expected voteshare is then modeled as a linear function of the average vote intention in the poll, plus some normally distributed random noise. Rothschild and Wolfers then back out a predicted value of the latent expected voteshare as a function of the estimated standard deviation of noise and the population proportion expecting the Democrat will win. As they note, this last step adds the assumption that the survey proportion is equivalent to the population proportion. Finally, the procedure itself requires variation in sample proportions, meaning that it is infeasible for use in a small number of polls, and does not allow researchers to test individual-level hypotheses regarding the determinants of predicted voteshares. Rather than going through this relatively cumbersome procedure, we believe a more fruitful approach is to simply ask respondents about their "latent" expected voteshare directly.

The coarseness of typical expectation measures also limits the use of predictions as dependent and independent variables. For example, studies of motivated reasoning might regress expectations on partisanship, finding that 90% of both parties believe their favored candidate will win. Without measuring expected voteshare, however, it is difficult to characterize the magnitude of motivated reasoning in this case. Did the 90% who expected the loser believe their candidate would win by a landslide, or by a few points? Were the 90% who expected the winner more or less "accurate" than those who predicted the loser? Aside from this issue of interpretation, the possibility of a 90% split on a binary variable greatly reduces the amount of variation in the dependent variable, mechanically preventing the discovery of meaningful differences. This lack of variation is also problematic in the context of elections as experiments: validating regression discontinuity designs via expectations requires precise measurements of anticipated outcomes, while showing election effects are strongest among surprised respondents is more difficult when surprise is measured with error.

4 Can Respondents Handle Numerical Predictions?

That existing measures of citizen predictions ask about who will win, but not by how much, is likely due to most survey researchers' avoidance of numerical questions. In a well-known example, citizens vastly overestimate the portion of the federal budget spent on foreign aid, which may reflect ignorance (Gilens 2001) or a lack of effort (Krosnick 1991). From findings such as this, survey researchers have apparently inferred that citizens are incapable of reporting numerical quantities with accuracy. Yet in a recent paper, Ansolabehere, Meredith, and Snowberg (2013) show this conclusion is unwarranted. At least in the domain of economic phenomena, these authors show that respondents can sensibly and accurately report numerical estimates of quantities such as gas prices and unemployment.

Does the ability to handle numerical questions extend to predictions of future political phenomena that are familiar to ordinary citizens? If so, then questions that ask about voteshare predictions may outperform existing measures of winner predictions. To explore this possibility, we conducted an original survey of approximately 2700 U.S. citizens between October 17 and 31, 2012, just prior to the 2012 presidential election.[1] The survey also included a second wave fielded between November 13 and 27, in which about half of the initial sample participated. Respondents were recruited online by Survey Sampling International (SSI), based on a constructed target population matching the census population of U.S. adults on age, gender, education, income, and geography, and were routed to the Qualtrics online survey platform.[2] The resulting responses come from a

[1]Replication data may be found in Sances and Quek (2015). We restrict the analysis to U.S. citizens aged 18 and above. Our population of interest is eligible voters in the United States, regardless of whether they actually turn out to vote. At no point in the analysis do we restrict the analysis to actual or likely voters.

[2]SSI recruits participants into its panel via online communities, social networks, and website ads, with specific efforts to incorporate less accessible demographic groups such as the elderly. SSI screens the participants before recruiting them into the panel, and randomly draws from the panel in its survey invitations. Some publications in political science using SSI samples include Kam (2012), Malhotra, Margalit, and Mo (2013), and Berinsky, Margolis, and Sances (2014). For further details on SSI, please see http://www.surveysampling.com. For more information on the Qualtrics platform,

https://doi.org/10.1093/pan/mpv022 Published online by Cambridge University Press

diverse national sample that is reasonably representative on these demographic observables (see Supplementary Materials). Because the SSI targeting strategy was successful in recruiting a representative sample, all of our analysis uses unweighted data.[3]

We measured respondents' expectations of the election outcome in two ways. The first method resembles the discrete measure of expectations used by existing studies. This question asked respondents:

> Thinking about the upcoming presidential election, do you think Mitt Romney or Barack Obama will win the election?

Respondents could then select one of the following five options:

> Mitt Romney definitely will be president
> Mitt Romney probably will be president
> Mitt Romney and Barack Obama have an equal chance of being president
> Barack Obama probably will be president
> Barack Obama definitely will be president

In most of the analyses below, we code respondents as predicting an Obama victory if they answered Obama would "probably" or "definitely" win, and zero otherwise; all our results are robust to other codings.

Our second measure of expectations asks about voteshares. This question reads:

> Thinking about only the votes cast for the two major parties, what percentage of the vote do you think Barack Obama and Mitt Romney will each receive in the NATIONAL VOTE?

Respondents were asked to enter their percentage point predictions for Barack Obama and Mitt Romney. The Obama prediction and Romney prediction were automatically summed by the computer and displayed to the respondent, and respondents were alerted and asked to revise their responses if their two predictions did not sum to 100%. Each response is coded as a quantitative voteshare prediction for Obama or Romney that ranges from 0% to 100%. As in Ansolabehere, Meredith, and Snowberg (2013), we drop respondents giving extreme values by omitting responses of 0% and 100%, though our results are robust to including these respondents.[4]

Before comparing responses to our two prediction measures, we first ask whether respondents can make accurate predictions about voteshare. Figure 1 shows four histograms of predicted voteshare among the full sample and by partisan subgroup. As the top-left panel of the figure shows, predictions among the full sample are approximately normally distributed, with a mass-point slightly above 50%. The overwhelming majority (83%) of predictions clusters within the 40%–65% range, and the 50%–55% range and 45%–50% range are the two highest-frequency categories. Among the full sample, the average prediction is only 1.1% higher than the actual

please see http://www.qualtrics.com/. Our module was part of a voting behavior study in which respondents were surveyed on the upcoming presidential election.

[3]In addition to predictions, we also measured a variety of factors that may affect the respondent's voteshare prediction. These include age, gender, race, education, household income, state of residence, political knowledge, and party identification. We measure knowledge as the sum of correct responses to four factual questions about politics. We measure party identification as an ordinal variable coded from 1 to 7 based on whether the respondent self-identifies as a Democratic or Republican partisan. The response categories range from Strong Democrat (1) to Strong Republican (7). Independents are those who do not self-identify as Democrat or Republican, and do not lean toward either party. Respondents who do not report themselves as Democrat or Republican but express a leaning toward either party are coded as a partisan of the party. This follows the finding that "leaning partisans" tend to be more partisan than those who report themselves as weak supporters of the Democratic or Republican party (Keith et al. 1992). The Supplementary Materials describe these variables in detail.

[4]We drop extreme predictions, as they represent either the failure of respondents to understand the question, and thus contribute random noise, or partisans who intentionally misrepresent their expectations to engage in "cheerleading" (Bullock et al. 2015). As we discuss in Section 6, this form of non-random measurement error could lead to bias in interaction effect estimates. Despite these issues, we replicate our main results in the Supplementary Materials when including the full range of estimates. We chose not to provide respondents with information about historical election results, which may have biased our results toward finding accuracy. Ansolabehere, Meredith, and Snowberg (2013) find that providing respondents with historical unemployment rates has little effect on average unemployment estimates, and "substantially reduces, but does not eliminate" unrealistically high responses (59).

https://doi.org/10.1093/pan/mpv022 Published online by Cambridge University Press

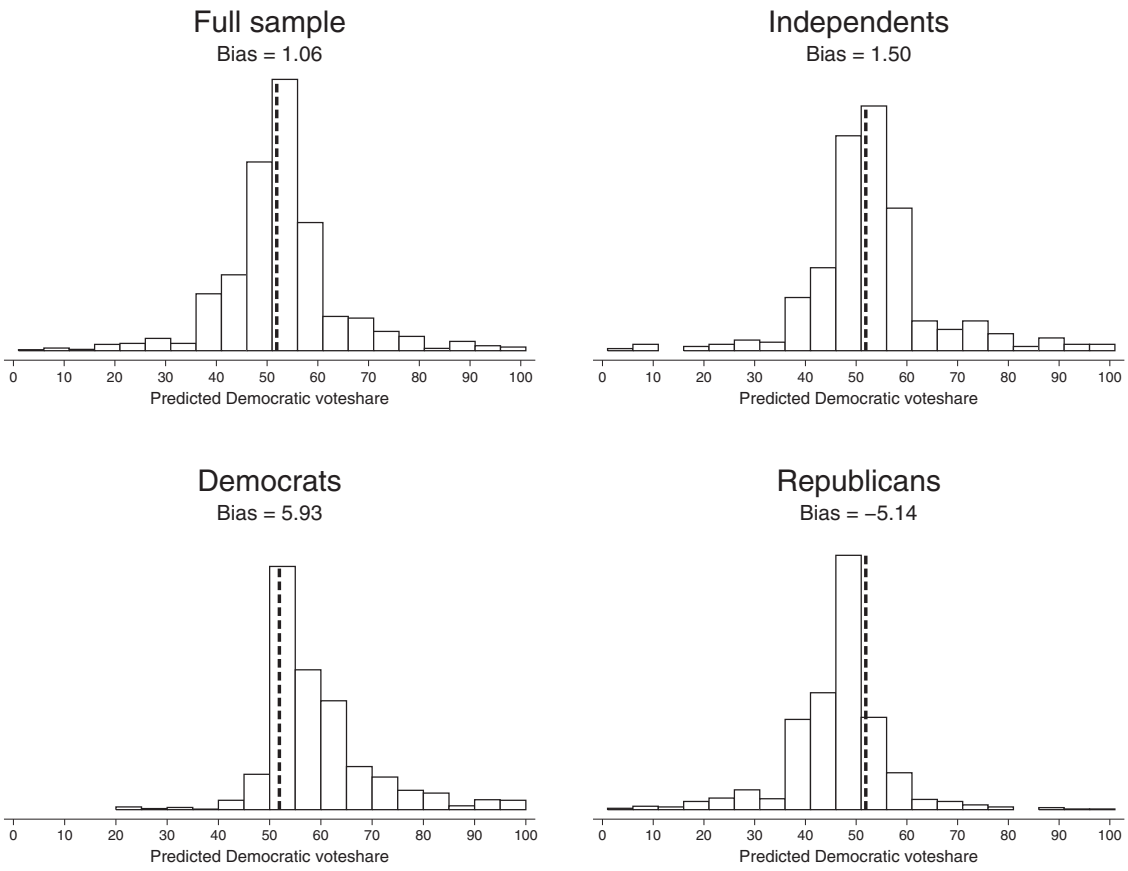


Fig. 1 Citizen forecasts of the 2012 election: Asking about voteshares.

outcome in the 2012 presidential election, indicated by the vertical dashed line at 51.9%.[5] By way of comparison, expert forecasts published in the October 2012 issue of *PS: Political Science and Politics* had an average error of 1.7% (Campbell 2012), and the October 31, 2012, Iowa Electronic Market forecasted an outcome of 50.5% for Obama, or 1.4% lower than the actual result.[6]

The following three panels show voteshare predictions for each partisan subgroup. In the top-right panel, we show the histogram for independents. Among partisan subsamples, independents are the most accurate predictors, with an average prediction that is 1.5 percentage points higher than the actual outcome. Predictably, identifiers with the major parties overestimated the voteshares for their preferred candidates. As shown in the bottom two histograms, Democrats overestimated Obama's voteshare by 5.9 points, whereas Republicans underestimated Obama's voteshare by 5.1 points. Note, however, that both Republican and Democrat respondents appeared to recognize the closeness of the race, with neither groups predicting a blowout for their candidate.

In Figure 2, we compare our voteshare prediction to two more conventional measures: vote intentions (the proportion of respondents saying they would vote for Obama) and winner

[5]We can confidently reject the null hypothesis that the average response is 50%, which would reflect random guessing ($p < 0.001$ for all four plots).

[6]See https://iemweb.biz.uiowa.edu/WebEx/marketinfo_english.cfm?Market_ID=350. We do not claim that our expected voteshare measure is generally superior to either expert or prediction market forecasts of voteshares. Indeed, because we only have an aggregate voteshare prediction for one survey and one election, a comprehensive comparison of voteshare predictions against the alternatives is beyond the scope of our study. Instead, we simply note that our measure performs relatively well in the context of the 2012 election, which supports our claim that respondents can produce sensible voteshare predictions.

https://doi.org/10.1093/pan/mpv022 Published online by Cambridge University Press

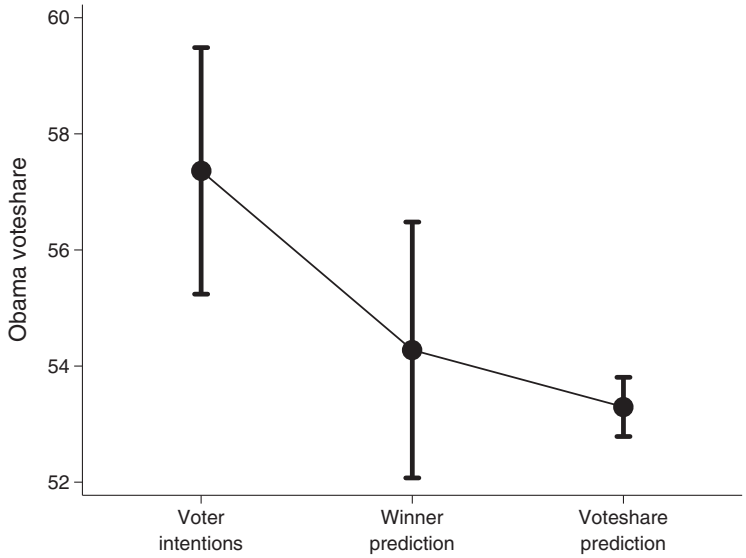


Fig. 2 Expectations are more accurate than intentions, but voteshare expectations are more accurate than winner expectations.

predictions (the proportion who believe Obama would win). As in previous work (e.g., Rothschild and Wolfers 2013), we see that forecasts based on intentions perform poorly relative to predictions: in our sample, about 57% of respondents reported planning on voting for the Democratic candidate, whereas about 54% of respondents predicted the Democrat would win. Thus, even this admittedly crude measure—the simple proportion of respondents who believed Obama would win—performs significantly better than the intention question. Transforming this proportion into a voteshare would likely yield a more accurate prediction, though this is not possible given we only have one sample, and also unnecessary given that we asked directly about voteshares. Indeed, we find that asking directly about voteshares performs the best: not only is the average prediction closest to the actual outcome, but the uncertainty around this estimate is dramatically smaller. The 95% confidence interval (calculated as 1.96 times the standard error of the mean) has a width of 4.2 and 4.4 percentage points for the vote intention and winner prediction questions, respectively, but only 1 percentage point for the voteshare prediction.[7]

5 Partisan Bias in Election Predictions: The Role of Information

Citizen predictions are not only potentially useful as forecasting tools, but they can also yield insight into the nature of partisan bias. However, commonly used measures of predictions may give misleading answers. To show this, Fig. 3 compares the performance of different prediction measures by partisan subgroup. Particularly striking is the reduction in bias among Democrats and Republicans. According to the winner prediction, roughly 80% of partisans believed that their candidate would win the election. We might therefore conclude that partisans experienced the 2012 campaign in radically different ways. The conclusion could subsequently sustain various plausible hypotheses—for instance, that partisans followed media outlets that overstated their favored candidate's support—that are, in fact, misleading. The voteshare predictions reveal a very different conclusion: using these measures, we see that both groups of partisans thought the election would be close, with their own candidate holding only a slight advantage over the other.

[7]The average voteshare prediction is slightly less accurate in Fig. 2 than Fig. 1, because in Fig. 2, we restrict the analysis to respondents who provided a valid answer to the vote choice question. Note also that the survey on which our prediction questions were placed randomly assigned some respondents to a non-standard vote intent question. These conditions are described in the Supplementary Materials; restricting the sample to the standard vote intent question, 56% of the sample intended to vote for Obama.

https://doi.org/10.1093/pan/mpv022 Published online by Cambridge University Press

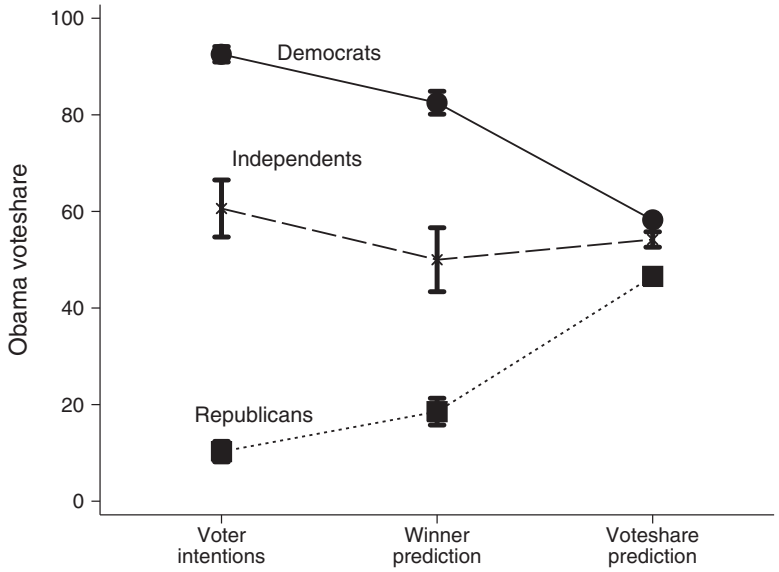


Fig. 3 Both vote intentions and winner predictions exhibit greater partisan bias than voteshare predictions.

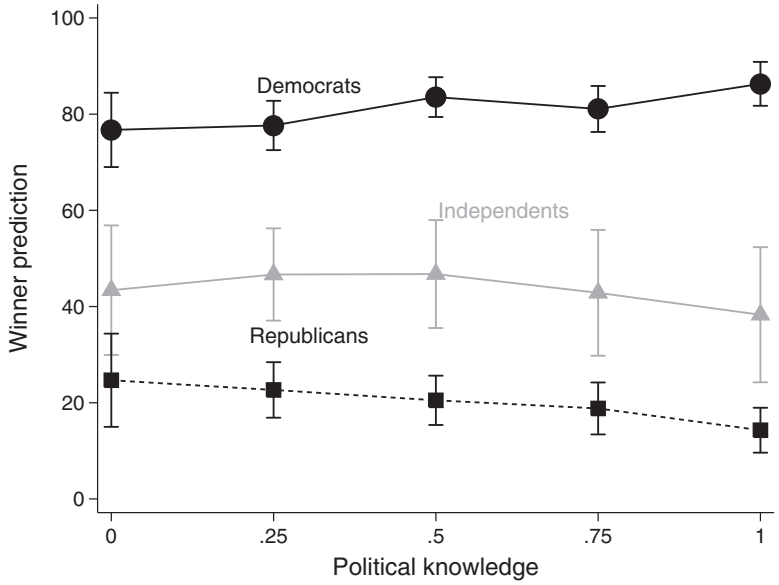


Fig. 4 Testing for motivated reasoning using winner predictions.

The increased precision of the voteshare prediction also leads to starkly different conclusions about potential solutions to partisan bias. According to theories of motivated reasoning, partisans will seek out information that favors the image of their own party, while avoiding information that harms this image (Taber and Lodge 2006). Thus, efforts to cure partisan bias through information will backfire, as partisan believers will only accept information that conforms to their pre-existing views, exacerbating perceptual biases.

Does this type of motivated reasoning occur in election predictions? This is not only important for learning about partisan bias, but also for the validity of election predictions as a forecasting tool, as such biases may skew aggregate forecasts. Once again, however, the answer to this question depends on how we measure predictions. In Figs. 4 and 5, we plot average predictions against factual knowledge about politics, with separate series for Democrats, Republicans, and independents. Figure 4, which uses the winner prediction on the vertical axis, leads one to believe that

https://doi.org/10.1093/pan/mpv022 Published online by Cambridge University Press

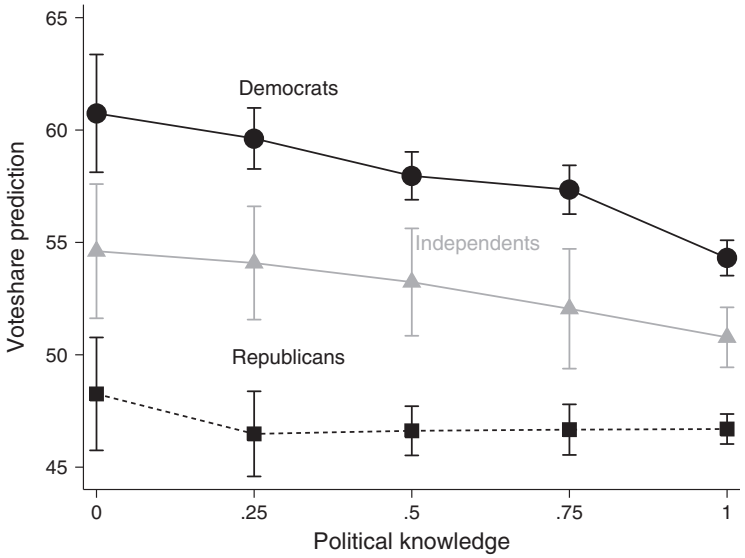


Fig. 5 Testing for motivated reasoning using voteshare predictions.

motivated reasoning is occurring. As political knowledge increases, the proportion of Democrats predicting a Democratic victory moves further from the truth. Republicans move symmetrically, predicting a smaller Democratic voteshare as knowledge increases.

Yet, Figure 5, which has voteshare predictions on the vertical axis, yields a very different conclusion. Now, we see that the average predicted voteshare among Democrats moves closer to the actual election outcome as knowledge increases. Likewise, the trend among Republicans is much less clear than before, and actually appears to flatten out after moving from the first to the second knowledge category. Democrats may have become more certain that Obama would win as their knowledge increased, but they were simultaneously realizing that his win margin would be smaller.

We next validate these graphical patterns using regression analysis. For both the winner prediction and the voteshare prediction, we regress predictions on party, political knowledge, and their interaction. In some specifications, we also include standard demographic measures, including education, income, intended voter turnout, age, gender, and race. For interpretation, we recode all variables to lie between zero and one, where a zero represents the minimum value and one the maximum. We exclude pure independents from this analysis, focusing only on major party identifiers as in Figs. 4 and 5. We show these results in Table 1.

As shown in the first row of column (1) of Table 1, the least informed Democrats and Republicans are about 50 percentage points apart in their expectations of who will win the election. As in the graphical analysis, information widens this gap: the coefficient on political knowledge indicates that a move from the minimum to the maximum knowledge among Republicans is associated with about a 10 point decrease in the expectation that Obama would win. The estimate in the third row indicates that Democrats' expectation that Obama would win also increases by 10 points as information increases.[8] These estimates are all statistically significant at the 0.05 level, and adding the control variables in column (2) leaves these results essentially unchanged. In contrast, the estimates using voteshare predictions show a narrowing of the gap with information. In column (3), the third row indicates that a change in knowledge significantly reduces Obama's predicted voteshare among Democrats by about 6 points, whereas the second row indicates Republicans are unaffected by information (the coefficient is −0.58 with a standard error of 1.00). Column (4), which adds controls to the analysis, reinforces this conclusion.

[8]While the coefficient on the interaction is 19 points, the marginal effect of information for Democrats is $\beta_{\text{Political knowledge}} + \beta_{\text{Dem X knowledge}} \approx -9 + 19 = 9$.

https://doi.org/10.1093/pan/mpv022 Published online by Cambridge University Press

Table 1 Voteshare predictions yield different inferences about partisan bias and information

	Winner prediction		*Voteshare prediction*	
	(1)	*(2)*	*(3)*	*(4)*
Democrat	51.17	48.84	14.58	13.13
	(3.76)	(3.87)	(1.08)	(1.08)
Political knowledge	–9.42	–9.23	–0.58	1.20
	(4.14)	(4.30)	(1.00)	(1.03)
Dem X knowledge	19.09	20.83	–6.30	–5.76
	(5.54)	(5.55)	(1.42)	(1.39)
Some college		3.18		–0.83
		(2.26)		(0.66)
Bachelors		0.72		–0.90
		(2.67)		(0.65)
Postgraduate		1.33		–1.24
		(2.82)		(0.64)
Income		4.08		–1.82
		(4.18)		(1.15)
Turnout		–2.39		–0.15
		(3.07)		(0.79)
Age		–7.93		–3.21
		(4.09)		(1.01)
Female		–2.49		0.52
		(1.85)		(0.45)
Black		10.28		5.76
		(3.03)		(0.96)
Other race		1.95		3.00
		(3.01)		(0.88)
Constant	25.09	27.85	47.13	48.82
	(2.86)	(4.70)	(0.80)	(1.28)
Observations	2032	2032	2032	2032

Note. Standard errors in parentheses.

These conflicting results likely reflect the fact that the different prediction measures capture different constructs once we move to the individual level. At the aggregate level, a voteshare prediction from the population that is close to 50% reflects aggregate uncertainty over whether Obama will win. The measure, however, does not necessarily capture uncertainty at the individual level: a prediction that Obama would win by 52–48 is not necessarily more uncertain than a prediction that Obama would win by 54–46.[9]

6 Enhancing Election-Based Natural Experiments with Voteshare Predictions

Besides forecasting elections and studying partisan bias, quantitative predictions are a useful tool when leveraging elections as natural experiments. In addition to discontinuity designs, numerous scholars have used panel methods which exploit unexpected election outcomes to learn about the effects of political control on public opinion (Ginsberg and Weissberg 1978; Finkel 1985; Keele

[9]In some applications, researchers may be more interested in certainty of victory versus predicted voteshare, though we argue this is not the case for motivated reasoning studies. Viewed in isolation, it may seem that a more certain victory for one's party makes one's party look better. Yet if partisans simultaneously believe their party will receive fewer votes, they are implicitly making their party look less electorally popular, a result that is hard to square with theories of motivated reasoning. In cases where researchers are interested in measuring uncertainty, it is likely that the binary prediction measure still leaves much to be desired. Instead of relying on this noisy measure, researchers could ask respondents about the probability of victory (as in Miller et al. 2012), or to estimate the highest and lowest possible voteshares each candidate could receive.

https://doi.org/10.1093/pan/mpv022 Published online by Cambridge University Press

2005; Anderson et al. 2005; Moehler 2009; Matsubayashi 2013; Sances and Stewart 2015; Margolis 2015).

To show how voteshare predictions can enhance these natural experiments, we focus on a particular type of election effect, namely the impact of elections on economic perceptions. Much research has documented that members of the president's party report feeling better about the economy than members of the opposite party (Bartels 2002). To test whether this pattern is the result of differing beliefs about the parties' relative competence, rather than selective perception or different evaluative criteria, Gerber and Huber (2010) exploit the 2006 midterm Congressional election, when the Democrats took control of the U.S. Congress, as an unexpected change in political control. Comparing pre- and post-election perceptions of the economy, Gerber and Huber find that Republican expectations decreased sharply relative to Democrats, a change they attribute to the change in political control resulting from the election.

Our own survey also included a measure of economic perceptions, as well as a pre- and post-election wave. As in Gerber and Huber (2010), our economic perception measure asked respondents:

> Thinking about the economy in the country as a whole, over the next year do you expect the nation's economy to get better, stay the same, or get worse?

Respondents then chose one of the following options:

> Get much better
> Get better
> Stay about the same
> Get worse
> Get much worse

We then use this measure, in combination with our voteshare prediction question, to test the validity of the 2012 election as a natural experiment. In particular, we test whether we can replicate previous findings regarding the effects of elections on economic perceptions, as well as whether these findings are stronger among those who were most surprised by the result.[10] Notably, while control of the White House did not change hands in 2012, most observers predicted the race would be close, leaving much room for surprised voters. Moreover, the potentially weaker treatment effect that results from the lack of change in control works against us finding an effect.

We first graph the basic results in Fig. 6. The vertical axis in this figure represents the change in economic perceptions across waves, with higher values indicating more positive shifts in economic perceptions; the horizontal axis represents the respondent's predicted voteshare for Barack Obama. We then plot the change in perceptions for Democrats (solid line) and Republicans (dashed line) using a nonparametric smoothing method, adding 95% confidence intervals for illustration.[11] Finally, the horizontal lines at about zero and –0.7 represent the unconditional average change for Democrats and Republicans, respectively.

The figure shows that we can replicate Gerber and Huber's finding in the 2012 election: Democrats were on average unaffected by the election outcome, whereas Republicans decreased their economic perceptions by an average of about 0.7 points on the five-point scale. More importantly, the conditional effect of party comports with an interpretation of the election as a natural experiment: among Republicans who predicted a larger Republican victory, economic perceptions decrease even more, about twice as much among the most surprised. In contrast, Democrats, who by and large predicted Obama would receive more than 50% of the vote, exhibit no change in their perceptions of the economy.

[10]We define surprise as the gap between expected voteshare and actual voteshare. Alternatively, researchers might separately measure respondents' uncertainty about the outcome, and test if effects are stronger among those who were more certain of their incorrect predictions. Again, we distinguish between expected voteshare and uncertainty here: Republicans who predicted Romney would receive 54% are more surprised than those who predicted he would receive 52%, but we do not claim that the latter are more or less certain in their prediction than the former.

[11]We use the "fpfit" command in Stata, which utilizes factional polynomial smoothing. For this figure, we drop predictions lower than the first percentile or higher than the ninety-ninth percentile.

https://doi.org/10.1093/pan/mpv022 Published online by Cambridge University Press

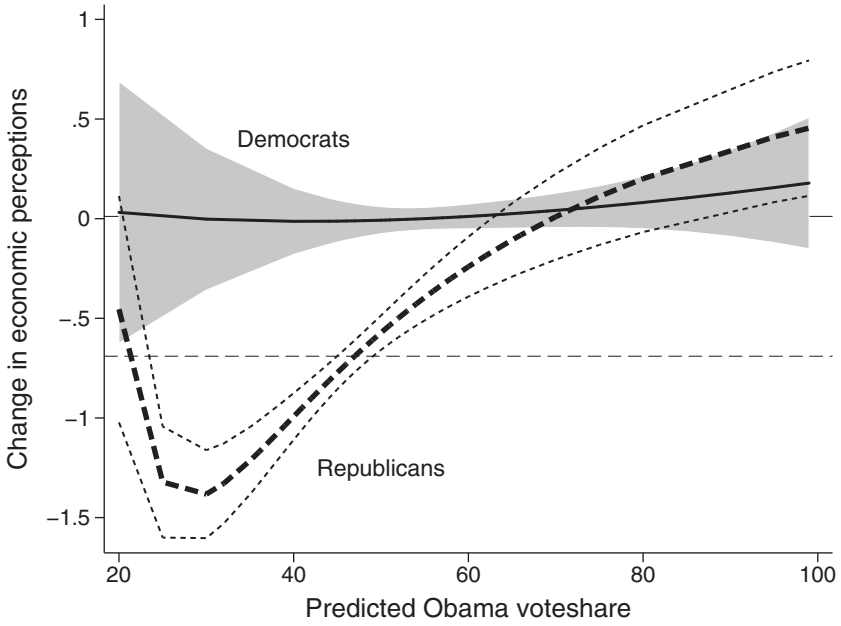


Fig. 6 Republicans who were more surprised at the election outcome changed their economic perceptions more.

Once again, we validate these patterns using multivariate analysis presented in Table 2. In the first two columns, we present the baseline effect by regressing the change in economic perceptions on an indicator for Republicans. With or without controls, we find that economic perceptions significantly decline by about 0.7 points for Republicans, relative to Democrats, as a result of the election. In the next four columns, we ask how the election effect varies with surprise. First, we ask how the effect varies with surprise measured using the respondent's prediction of who will win. Using this simple bivariate measure does help to validate the experiment, as we see that the election effect is concentrated among those who believed that Obama would lose: the coefficient on Republican in columns (3) and (4) suggests that those who believed Obama would lose decreased their economic perceptions by about 0.8 points. Further, the interaction between party and predictions indicates that this decline shrinks to about $(-0.81 + 0.56 \approx)$ 0.3 points among those who anticipated Obama's re-election.

While showing that effects are concentrated among surprised respondents is a useful validity check, we may also want to use predictions to detect treatment effects among particular subsamples, such as those most surprised by the outcome. Unfortunately, the crudeness of the binary prediction measure could potentially deflate the true effect of political events. This is particularly problematic in a case such as ours, where the absence of a shift in political control could make it harder to detect an effect. To illustrate this problem, and to show how voteshare predictions provide a solution, in columns (5) and (6) we use our voteshare prediction as the key interactive variable. Using this measure reveals that the election effect is not only concentrated among surprised respondents, but it is substantially strengthened. Among Republicans who predicted the largest Romney victory, the election effect is about 1.5 points (standard error = 0.3), which is over twice as high as the unconditional average of 0.7 points. Likewise, the interactive coefficient is over three times as large as when using the binary measure, between 1.8 and 2 points (standard error = 0.6) depending on the specification.

In addition to highlighting advantages, we also offer a note of caution for studies of electoral effects that seek to leverage election predictions. As with any analysis of effect heterogeneity, it is important to consider whether the moderating variable, in this case predicted vote margins, is plausibly unrelated to other factors that may also affect the outcome, in this case economic perceptions. One violation of this assumption would be if those who expected a Romney victory are stronger partisans, who may also be more likely to change their economic perceptions as a result of a Romney loss. Such an effect could arise as a result of "partisan cheerleading," which could cause

https://doi.org/10.1093/pan/mpv022 Published online by Cambridge University Press

Table 2 Voteshare predictions are better able to detect election effects

	Baseline		*X winner prediction*		*X voteshare prediction*	
	(1)	*(2)*	*(3)*	*(4)*	*(5)*	*(6)*
Republican	–0.72	–0.69	–0.85	–0.80	–1.53	–1.59
	(0.05)	(0.05)	(0.09)	(0.09)	(0.30)	(0.30)
Prediction			–0.03	–0.01	0.35	0.18
			(0.08)	(0.08)	(0.31)	(0.33)
Rep X prediction			0.57	0.56	1.84	1.99
			(0.12)	(0.12)	(0.58)	(0.58)
Political knowledge		–0.20		–0.18		–0.19
		(0.09)		(0.08)		(0.09)
Some college		0.06		0.04		0.06
		(0.07)		(0.07)		(0.07)
Bachelors		0.05		0.04		0.05
		(0.07)		(0.07)		(0.07)
Postgraduate		0.25		0.24		0.24
		(0.07)		(0.07)		(0.07)
Income		–0.02		–0.03		–0.02
		(0.01)		(0.01)		(0.01)
Turnout		–0.03		–0.02		–0.02
		(0.08)		(0.07)		(0.08)
Age		–0.00		–0.00		–0.00
		(0.00)		(0.00)		(0.00)
Female		–0.01		–0.00		–0.00
		(0.05)		(0.05)		(0.05)
Black		0.01		0.01		0.01
		(0.08)		(0.08)		(0.09)
Other race		0.03		0.03		0.02
		(0.10)		(0.10)		(0.09)
Constant	0.01	0.45	0.04	0.44	–0.19	0.32
	(0.03)	(0.14)	(0.07)	(0.15)	(0.18)	(0.22)
Observations	1175	1175	1175	1175	1175	1175

Note. Standard errors in parentheses.

respondents to misreport their beliefs about election outcomes and the economy in a way that makes their own party look better (Bullock et al. 2015, forthcoming).

While the severity of this issue ultimately depends on the application, we offer three responses to this concern here. First, predictions can be correlated with cheerleading only to the extent that it actually occurs, which is still an open question in the literature (Prior 2007; Gerber and Huber 2009; Bullock et al. 2015, forthcoming). Second, we indirectly tested this explanation by examining whether the effect of surprise varies with strength of partisanship. Given that cheerleading should be increasing with the strength of partisanship, the effect of surprise should be greater as the partisan strength increases. We find this is not the case.[12] Third, and as shown in Fig. 3 above, the quantitative prediction is actually less prone to partisan influence, and therefore less likely to be driven by cheerleading, than the qualitative prediction. If cheerleading were driving the effects of surprise, we should then observe a stronger effect when using the binary surprise measure; yet, we

[12]Specifically, we conducted a regression interacting six indicators for partisanship (one for each scale point) with the prediction variable, to test whether the positive relationship between perceptions and predicted voteshare holds among both strong and weak Republicans. We find that it does: the interactions between predicted voteshare and leaning, weak, and strong Republicans range between 1.5 and 2.1, and are always significant at the 0.10 level. A test of equality of these interactive coefficients yields an F-statistic of 0.21 ($p = 0.81$), which means we cannot reject the null hypothesis that the effects are equal regardless of partisan strength. Performing this test using the winner expectation yields an F-statistic of 0.49 ($p = 0.61$).

https://doi.org/10.1093/pan/mpv022 Published online by Cambridge University Press

find the opposite. Of course, this does not preclude the possibility that prediction measures will be confounded in other applications, and researchers should conduct their own tests of this threat as appropriate. As a preventive measure, researchers would do well to rely on quantitative prediction measures, which we have demonstrated to be less prone to partisan influence relative to qualitative measures.

7 Conclusion

Whether and how ordinary citizens can predict election outcomes has consequences over election forecasts, the study of partisan bias, and the use of elections as natural experiments. Existing measures of prediction, however, overwhelmingly ask people to predict the victor rather than the voteshare. The conventional approach seemingly assumes a qualitative measure is preferable, as ordinary citizens are unwilling or unable to process quantitative predictions.

Suspending that assumption, we designed and fielded original nationwide surveys around the 2012 U.S. presidential election to test whether and how citizens can predict election outcomes. Our research reveals that ordinary citizens can make sensible voteshare estimates at a level of accuracy that rivals that of professional forecasters. After showing that respondents can handle numerical predictions, we uncover how traditional qualitative forecasts can lead to severely mistaken inferences. Specifically, an investigation of partisan bias shows that qualitative predictions vastly overstate the degree of partisan bias in election forecasts and lead to wrong conclusions, such as how political knowledge exacerbates bias when, in fact, it does not. To demonstrate the utility of quantitative predictions in election-based natural experiments, we also use the 2012 election as a natural experiment to test for an effect of partisan control on economic perceptions. We show that the effect of the election increases further as the respondent predictions strayed from the true voteshare, replicating a recent result (Gerber and Huber 2010), and validating both our prediction measure and the natural experiment itself.

Future work may extend our research in several directions. For example, to what extent can voteshare predictions be combined with other methods to further increase the precision of election predictions? One interesting possibility is to connect our measure with Barber et al.'s (2014) method to identify a more precise sample of the likely electorate. Such a strategy may not only increase precision, but could also allow for better testing of the relationship between turnout propensity and perceived electoral closeness. A second avenue for future research is to investigate precisely how respondents form expectations about election outcomes. For instance, do respondents form voteshare expectations based on social networks (e.g., Rothschild and Wolfers 2013)? Alternatively, are more accurate predictors simply more likely to hear elite predictions in the news media? Our results for pure independents in Fig. 2 seem to cast doubt on the latter explanation: even the least informed independents appear to do as well as or better than the most informed partisans. Given our study contains only a small number of uninformed independents and lacks any measure of respondent networks, however, we must ultimately leave the question of mechanisms to future work.

Practically speaking, our findings should interest methodologists and pollsters, as well as political scientists and interdisciplinary scholars of collective intelligence. Together, the results show the inferential drawbacks in existing measures of election prediction; challenge the assumption that ordinary people cannot process quantitative predictions; and demonstrate the benefits of a simple but powerful remedy: to ask people to make direct numerical voteshare predictions (if one trusts the respondents more), or at least, to ask people to predict both the victor and the voteshare (if one trusts the respondents less). Since the latter presents a dominant strategy, we advocate that the present practise of polling respondents on the victor but not the voteshare be discontinued.

More generally, our research—together with other recent research (Ansolabehere, Meredith, and Snowberg 2013)—speaks to the need to re-examine the assumptions that lurk behind the use of binary variables to capture variations that are continuous in nature. Even if the practice is conventional and appears innocuous, it has a set of natural drawbacks: it is less informative, because only binary information is extracted; it is less direct, because assumptions are required to reconvert the information into the true quantity of interest; and it is also more likely to foster inferential

https://doi.org/10.1093/pan/mpv022 Published online by Cambridge University Press

mischief, because of the reduction in variation imposed by the truncated measure. In the end, the assumptions underlying the use of a specific binary variable may or may not be justifiable, but their inferential costs should always be carefully considered.

References

Ansolabehere, Stephen, Marc Meredith, and Erik Snowberg. 2013. Asking about numbers: Why and how. *Political Analysis* 21:48–69.

Barber, Michael J., Christopher B. Mann, J. Quin Monson, and Kelly D. Patterson. 2014. Online polls and registration-based sampling: A new method for pre-election polling. *Political Analysis* 22:321–35.

Bartels, Larry M. 2002. Beyond the running tally: Partisan bias in political perceptions. *Political Behavior* 24:117–50.

Bullock, John G., Alan S. Gerber, Seth J. Hill, and Gregory A. Huber. 2015. Partisan bias in factual beliefs about politics. *Quarterly Journal of Political Science*, forthcoming.

Campbell, Angus, Philip E. Converse, Warren E. Miller, and Donald E. Stokes. 1960. *The American Voter*. New York: John Wiley and Sons.

Campbell, James E. 2012. Forecasting the 2012 American national elections. *PS: Political Science & Politics* 45:610–3.

Caughey, Devin, and Jasjeet S. Sekhon. 2011. Elections and the regression discontinuity design: Lessons from close US House races, 1942–2008. *Political Analysis* 19:385–408.

Daniller, Andrew M., Laura Silver, and Devra Coren Moehler. 2013. Calling it wrong: Partisan media effects on electoral expectations and institutional trust. Paper presented at the 2013 American Political Science Association Annual Meeting. http://ssrn.com/abstract=2301154 (accessed September 17, 2015).

Eggers, Andrew C., Anthony Fowler, Jens Hainmueller, Andrew B. Hall, and James M. Snyder. 2015. On the validity of the regression discontinuity design for estimating electoral effects: New evidence from over 40,000 close races. *American Journal of Political Science* 59:259–74.

Erikson, Robert S., and Christopher Wlezien. 2008. Are political markets really superior to polls as election predictors? *Public Opinion Quarterly* 72:190–215.

———. 2012. Markets vs. polls as election predictors: An historical assessment. *Electoral Studies* 31:532–9.

Finkel, Steven E. 1985. Reciprocal effects of participation and political efficacy: A panel analysis. *American Journal of Political Science* 29:891–913.

Gerber, Alan S., and Gregory A. Huber. 2009. Partisanship and economic behavior: Do partisan differences in economic forecasts predict real economic behavior? *American Political Science Review* 103:407–26.

———. 2010. Partisanship, political control, and economic assessments. *American Journal of Political Science* 54:153–73.

Gilens, Martin. 2001. Political ignorance and collective policy preferences. *American Political Science Review* 95:379–95.

Ginsberg, Benjamin, and Robert Weissberg. 1978. Elections and the mobilization of popular support. *American Journal of Political Science* 22:31–55.

Graefe, Andreas. 2014. Accuracy of vote expectation surveys in forecasting elections. *Public Opinion Quarterly* 78:204–32.

Herron, Michael C. 2000. Estimating the economic impact of political party competition in the 1992 British election. *American Journal of Political Science* 44:326–37.

Kam, Cindy D. 2012. Risk attitudes and political participation. *American Journal of Political Science* 56:817–36.

Keele, Luke. 2005. The authorities really do matter: Party control and trust in government. *Journal of Politics* 67:873–86.

Keith, Bruce E., David B. Magleby, Candice J. Nelson, Elizabeth Orr, and Mark C. Westlye. 1992. *The myth of the independent voter*. Berkeley, CA: University of California Press.

Knight, Brian. 2006. Are policy platforms capitalized into equity prices? Evidence from the Bush/Gore 2000 presidential election. *Journal of Public Economics* 90:751–73.

Krosnick, Jon A. 1991. Response strategies for coping with the cognitive demands of attitude measures in surveys. *Applied Cognitive Psychology* 5:213–36.

Lewis-Beck, Michael S., and Tom W. Rice. 1992. *Forecasting elections*. Washington, D.C.: CQ Press.

Lewis-Beck, Michael S., and Charles Tien. 1999. Voters as forecasters: A micromodel of election prediction. *International Journal of Forecasting* 15:175–84.

Malhotra, Neil, Yotam Margalit, and Cecilia Mo. 2013. Economic explanations for opposition to immigration: Distinguishing between prevalence and conditional impact. *American Journal of Political Science* 57:391–410.

Margolis, Michele M. 2015. *Don't lose control: How partisanship and the political landscape shape religious behaviors*. Working paper, Department of Political Science, University of Pennsylvania.

Matsubayashi, Tetsuya. 2013. Do politicians shape public opinion? *British Journal of Political Science* 43:451–78.

Miller, Michael K., Guanchun Wang, Sanjeev R. Kulkarni, H. Vincent Poor, and Daniel N. Osherson. 2012. Citizen forecasts of the 2008 US presidential election. *Politics & Policy* 40:1019–52.

Moehler, Devra C. 2009. Critical citizens and submissive subjects: Election losers and winners in Africa. *British Journal of Political Science* 39:345–66.

Murr, Andreas Erwin. 2011. Wisdom of crowds? A decentralised election forecasting model that uses citizens' local expectations. *Electoral Studies* 30:771–83.

Page, Benjamin I., and Robert Y. Shapiro. 1992. *The rational public: Fifty years of trends in Americans' policy preferences*. Chicago: University of Chicago Press.

https://doi.org/10.1093/pan/mpv022 Published online by Cambridge University Press

Prior, Markus. 2007. Is partisan bias in perception of objective conditions real? Paper presented at the Annual Conference of the Midwest Political Science Association in Chicago, Illinois.

Rothschild, David M. 2009. Forecasting elections comparing prediction markets, polls, and their biases. *Public Opinion Quarterly* 73:895–916.

Rothschild, David M., and Justin Wolfers. 2013. Forecasting elections: Voter intentions versus expectations. Working paper, Department of Economics, University of Michigan. http://users.nber.org/~jwolfers/Papers/VoterExpectations.pdf (accessed September 17, 2015).

Sances, Michael, and Kai Quek. 2015. Replication data for: Closeness counts: Increasing precision and reducing errors in mass election predictions, V1 [Version]. http://dx.doi.org/10.7910/DVN/INGQYA (accessed July 20, 2015).

Sances, Michael, and Charles Stewart. 2015. *Partisanship and confidence in the vote count: Evidence from U.S. National Elections since 2000*. Working paper, Department of Political Science, Vanderbilt University.

Snowberg, Erik, Justin Wolfers, and Eric Zitzewitz. 2007a. Partisan impacts on the economy: Evidence from prediction markets and close elections. *Quarterly Journal of Economics* 122:807–29.

———. 2007b. Party influence in Congress and the economy. *Quarterly Journal of Political Science* 2:277–86.

Taber, Charles S., and Milton Lodge. 2006. Motivated skepticism in the evaluation of political beliefs. *American Journal of Political Science* 50:755–69.

Thibodeau, Paul, Matthew M. Peebles, Daniel J. Grodner, and Frank H. Durgin. 2015. The wished-for always wins until the winner was inevitable all along: Motivated reasoning and belief bias regulate emotion during elections. *Political Psychology*, 36: 431–48.

Wolfers, Justin, and Eric Zitzewitz. 2004. Prediction markets. *Journal of Economic Perspectives* 18:107–26.

https://doi.org/10.1093/pan/mpv022 Published online by Cambridge University Press