Are Polls and Probabilities Self-Fulfilling Prophecies?

WORD COUNT: 2749

David Rothschild
Microsoft Research
David@ResearchDMR.com
www.ResearchDMR.com

Neil Malhotra
Stanford University
malhotra_neil@gsb.stanford.edu
http://www.stanford.edu/~neilm/

This draft: March 26, 2013

Abstract

Psychologists have long observed that people conform to majority opinion, a phenomenon sometimes referred to as the "bandwagon effect." In the political domain people learn about prevailing public opinion via ubiquitous polls, which may produce a bandwagon effect. Newer types of information—published probabilities derived from prediction market contract prices and aggregated polls—may play a similar role. Consequently, polls and probabilities can become self-fulfilling prophecies whereby majorities, whether in support of candidates or policies, grow in a cascading manner. Despite increased attention to whether the measurement of public opinion can itself affect public opinion, the existing empirical literature is surprisingly limited on the bandwagon effects of polls and non-existent on the bandwagon effects of probabilities. To address this gap, we conducted an experiment on a diverse national sample in which we randomly assigned people to receive information about different levels of support (or probability of passage) for three public policies. We find that public opinion as expressed through polls significantly impacts individual-level attitudes whereas probabilities exhibit no effect. We also posit a mechanism underlying the bandwagon effect for polls: low public support decreases support for policies but high public support does not increase support.

Psychologists have long observed that people conform to majority opinion (e.g. Sherif 1936; Asch 1951; Deutsch and Gerard 1955). This is especially possible in the domain of politics where people regularly learn the views of the majority via public opinion polling. During elections, much of the media coverage focuses on the "horse race," or the levels and changes of candidate support (e.g. Broh 1980; Brady and Johnston 1987; Mutz 1995). Outside of election coverage, the media often report polls about public policy questions. The bandwagon effect occurs when people change their opinions to conform to the majority, shifting their preferences in favor of the leading candidate or the most popular policy position (Simon 1954). Bandwagon effects can make polls self-fulfilling prophecies; a poll's prediction may come to pass not only because it measures public opinion but also because it may influence public opinion.[1]

While polling has been a central aspect of politics for decades, new innovations such as prediction markets and forecasts based on aggregated polling have led to an increased opportunity for the public to learn about the *probability* of a political outcome occurring.[2] In prediction markets such as Betfair, Intrade, and IEM, people can purchase securities that pay $1.00 if a certain election outcome occurs (e.g. "Barack Obama reelected president of the United States"). Assuming efficient markets, the price of a contract is equivalent to the probability of the event occurring (Wolfers and Zitzewitz 2007; Manski 2004). Along with prediction markets, the public is often exposed to probabilities derived from aggregating polls (e.g. Nate Silver's FiveThirtyEight). By aggregating and de-biasing individual polls into single summaries of probabilities, these forecasts are efficient means for citizens to learn about public opinion.

[1] Note that the bandwagon effect, a form of conformity, is the mirror image of the false consensus effect, where people misperceive that their own behaviors and attitudes are more popular than they actually are (Ross 1977). In the political domain, one mechanism underlying the false consensus effect is wishful thinking, or people gaining utility from thinking their candidate is ahead or their opinions are popular (Granberg and Brent 1983; Krizan et al. 2010). Observational data cannot disentangle the bandwagon effect from these other phenomena because of simultaneity; an experimental approach is therefore required.

[2] Below we refer to information from public opinion surveys as *polling* and information about the probability of an outcome occurring from prediction markets or forecasts based on aggregated polling data as *probabilities*.

Many have recently raised concerns that by reporting public opinion, polls and probabilities can change individual-level attitudes. During the 2012 elections, many conservative commentators complained about "skewed" polls and how the media was attempting to influence the electorate through polling showing Barack Obama leading Mitt Romney (Easley 2012). Nate Silver himself has considered stopping releasing probabilities because of his fear that they are influencing attitudes and hindering the democratic process (Byers 2013). Despite increased media attention to the potential effects of polling on public attitudes, only a limited number of studies have explored bandwagon effects in the political domain.[3] Much of the extant data are quite dated and few studies use national samples.[4] Further, no study has explored how probabilities influence individual-level attitudes.

Social psychological research suggests three principle mechanisms by which polls may induce conformity: (1) normative social influence, or people's desire to adopt the majority position in order to feel liked and accepted or believe they are on the winning team (Deutsch and Gerard 1955); (2) informational social influence, or people learning from the "wisdom of crowds" via social proof because they "believe that others' interpretation of an ambiguous situation is more accurate than ours and will help us choose an appropriate course of action" (Aronson et al. 2005); and (3) people resolving cognitive dissonance by switching to the side they infer is going to win based on the poll (Kay et al. 2002). These mechanisms could also explain why probabilities could influence public opinion, with cognitive dissonance reduction perhaps the most theoretically sensible pathway.

We designed and conducted an experiment to assess whether polls and probabilities

[3] Studies that have explored the bandwagon effect include: Ansolabehere and Iyengar (1994); Kay et al. (2002); Mutz (1997. 1999); Marsh (1984); Lavrakas et al. (1991); Dizney and Roskens (1962); Navazio (1977); Ceci and Kain (1982); Lang and Lang (1984); Sinclair and Plott (2012); Rickershauser and Aldrich (2007); Goidel and Shields (1994); Mehrabian (1998); Morwitz and Pluinski (1996); McAllister and Studlar (1991).
[4] Exceptions are Lavrakas et al. (1991) and Mutz (1997, 1999).

affect public opinion on public policy issues. In addition to providing timely evidence on a topic that has received substantial attention, this research is innovative in several respects. First, we leverage a national, diverse sample, allowing us to reach more externally valid conclusions than research that has relied on local convenience samples. Second, we test the effects of probabilities alongside polls, allowing us to directly compare them. Third, our treatments cover the entire spectrum of support, meaning that we can test whether the underlying mechanism of the bandwagon effect is about the mobilizing power of public support or the demobilizing effects of public opposition. Fourth, while most studies of bandwagon effects have examined electoral contests, we explore the effects of polling and predictions on public policy attitudes (see Marsh 1984 for an exception).

We focus on preferences towards public policies because the role of polling in opinion cascades may help us understand how previously unpopular issues have increased in popularity (e.g. gay marriage; Pew 2012a) while previously popular issues have decreased in popularity (e.g. capital punishment; Pew 2012b). For instance, support for gay marriage has increased from 37% to 58% in less than ten years, a change that cannot solely be explained by cohort replacement (Langer 2013). Perhaps some of this increase in support is the result of bandwagon effects. Nonetheless, our findings speak to the literature on how voters incorporate information from sequential elections such as presidential primaries (e.g. Callander 2007; Bartels 1988; Knight and Schiff 2010; Morton and Williams 2001) or from early pre-election polling (e.g. Sinclair and Plott 2012; Rickershauser and Aldrich 2007).

We find strong evidence of a bandwagon effect of polls; people are much more supportive of policies that have higher general public support. The mechanism underlying this effect is that showing lower support demobilizes citizens; high support is not mobilizing.

Conversely, probabilities have basically no effect on opinion.

Experimental Design

Overview

We first provide a general overview of the experimental design before discussing specifics of how the outcomes were measured and the treatments were administered. Before any treatment information was presented, we measured support for three policies. We then asked some filler questions in order to provide some separation between the treatment and the measurement of the dependent variable; this protects against both consistency bias and stickiness of pre- and post-treatment responses.[5] The experiment has two levels of treatments. First, respondents were randomly assigned to receive one of two types of information: (1) a *poll* of the general population showing the level of support for a policy; and (2) a prediction about the *probability* of the policy passing. Respondents were then randomly assigned to a continuous variable representing the level of support for the target issue. We then measured support for the target policy again post-treatment.

Data

Data were collected as part of the 2011 Cooperative Congressional Election Study (CCES).[6] Interviews were conducted using an opt-in sample of 479 respondents over the Internet

[5] We asked fifteen filler questions which spanned five topics: abortion, Hurricane Katrina, unemployment, the Electoral College, and the minimum wage. For each topic, we asked a factual knowledge item and some follow-up questions about respondents' confidence in their answer.

[6] The CCES is a large-scale, omnibus survey on political issues which first asked a series of common content questions shared across all researchers adding questions to the platform, followed by the individual researchers' studies. The CCES has been conducted six times since 2006 in advance of the November elections and has featured the participation of over 50 research institutions. The target population was U.S. adult citizens. All interviews were conducted in English. Because respondents were members of an opt-in Internet panel, a standard response rate cannot be calculated. 41,436 panelists were invited to take the CCES survey, of which 11,355 completed the interview. After agreeing to take the survey, 479 respondents were randomly assigned to the module on which our questions appeared. Consequently, in expectation, the completion rate for our module is equal to the completion rate for the entire CCES survey. Item non-response was not an issue in our data as all 479 respondents completed the full set of questions used in this analysis. Except for one question about whether Afghanistan was a mistake in the

by YouGov/Polimetrix between November 9, 2011, and January 2, 2012. While there are concerns that respondents of opt-in Internet surveys are more politically interested, YouGov/Polimetrix has developed sampling techniques to mitigate these concerns.[7] Both layers of randomization were successful (see Appendix 1).

Dependent Variable

We asked respondents to report their level of support for three different policy proposals: reducing troop levels in Afghanistan, free trade agreements, and public financing of elections. We chose three topics that covered different aspects of American politics (foreign policy, economic policy, and election administration).[8] Because the dependent variable was asked both pre-treatment and post-treatment, we can control for pre-treatment attitudes to assess how much the treatments change policy support relative to the baseline of initial support. Question wordings for the three issues are presented in Appendix 2. The three issues were presented in a random order to respondents.

We asked respondents to report the probability that they would vote for the policy in a national referendum. Accordingly, the dependent variable ranges from 0 to 100. We employed this measurement strategy for two reasons, one substantive and one methodological. Substantively, the measure is intended to tap not only attitudes but also have a behavioral element (i.e., a vote intention). Methodologically, using a finely grained scale increases the

---

middle of the common content, there was no substantive overlap between the common content questionnaire and our own items.

[7] YouGov/Polimetrix uses a technique called sample matching to draw representative samples of the U.S. population from its panel of voluntary survey participants (Rivers n.d.). YouGov/Polimetrix draws samples from nationally representative probability surveys (e.g. the American Community Survey, the Current Population Survey, the Pew U.S. Religious Landscape Survey) and then matches panelists to this target sample based on observable characteristics such as age, race, education, technology usage, and several other factors. Recent studies have shown that YouGov/Polimetrix samples do as well as more traditional data collection techniques (e.g. RDD telephone interviewing) at matching known population benchmarks (Vavreck and Rivers 2008; Ansolabehere and Schaffner 2011).

[8] Below we refer to the dependent variable as *individual-level policy support* to distinguish it from the exogenously-provided public support for the policies presented in the treatment information.

opportunity for the experimental treatments to shift opinions. We conducted a follow-up study which shows that our measurement strategy taps attitudes in a similar way to a Likert scale.[9] The pre- and post-treatment questions are exactly the same, but we use a slightly different graphical interface pre- and post-treatment.[10] Distributions of pre- and post-treatment policy support are presented in Appendix 4.[11]

Treatments

For all three policy issues, respondents were provided with one of two types of treatment information: polls or probabilities. For both treatment types respondents randomly received a value ranging from 20 and 80 (inclusive) in increments of five. We refer to this value as the *treatment value*. The treatment values for the three issues were presented in a randomized order. To enhance realism, the three treatment values were different for each of the three issues.[12]

Respondents assigned to the poll condition read descriptions of the following form:

Below is the percentage of Americans who support [a meaningful reduction in U.S. troops in Afghanistan by June 30, 2012/more free trade agreements with North, Central, and South American countries/public financing of elections]. *This value is created by aggregating the best available polls.*

Respondents assigned to the probability condition read descriptions of the following form:

Below is the likelihood of there being a meaningful reduction of U.S. troops from

[9] The follow-up was conducted using 500 volunteers on Amazon's Mechanical Turk between January 25, 2013 and January 30, 2013. For all three issues we found extremely high Pearson correlations between the 100-point measure and five-point Likert scales ranging from “strongly disagree” to “strongly agree” (Afghanistan: $r = 0.72$; free trade: $r = 0.85$, public financing: $r = 0.90$). Further, the average value of the continuous vote intention monotonically maps with the rating scale categories (see Online Appendix 1).
[10] We measure pre-treatment policy support with a horizontal thermometer ranging from 0-100 whereas we measure post-treatment policy support with a vertical thermometer with the same range (see Appendix 3 for the graphical presentations). We ask the questions using slightly different visual presentations to avoid stickiness in responses while limiting confounding issues by not changing the wording of the questions.
[11] The variables are normally distributed with low density in the tails, suggesting that a few respondents dramatically shifting their positions do not drive the results. In fact, the modal respondent who does not update at all in response to the treatment information is one with extreme views, which is precisely what we would expect.
[12] At the end of the experiment, respondents were debriefed about the true levels for both treatment types based on current data.

Afghanistan by June 30, 2012. The U.S. currently has 100,000 troops in Afghanistan and a meaningful reduction is defined as 80,000 or less troops left. *This value is created by aggregating the best available forecasts.*[13]

We did not include a pure control group because the comparative static we are interested in testing is how individual policy support changes as the treatment value changes. The relevant baseline is not the absence of public opinion information. Further, the pre-treatment measures of the dependent variable can be used to assess policy support in the absence of any polls or probabilities.

Estimation Strategy

To test whether the value of the polling or probability affects individual-level policy attitudes, we estimate the following random effects regression model via OLS, pooling responses from all three issues together:

$$Y_{ij} = \beta_1 P_{ij} + \beta_2 X_{ij} + \alpha_j + \gamma_i + \varepsilon_{ij} \tag{1}$$

where i indexes respondents, j indexes issues, Y_{ij} represents support for issue j measured post-treatment for each respondent, P_{ij} represents support for policy j measured pre-treatment for each respondent, X_{ij} represents the randomized treatment value for the issues ranging from 20-80%, α_j represents issue dummies where free trade is the omitted category, γ_i is a random coefficient for each respondent assumed to come from a Gaussian distribution, and ε_{ij} represents stochastic error. Standard errors are clustered by respondent. As a robustness check, we also report results from a model including fixed effects for respondents. The coefficient of interest from model (1) is β_2, or the effect of increasing the treatment value by 1 percentage point on individual-level policy support. For example, a coefficient of 0.1 means that a ten-point increase in the treatment value increases policy support by 1 percentage point. Therefore, moving across the full range of

[13] The treatment information for all three issues is presented in Appendix 2.

possible treatment values (20 to 80) increases policy support by 6 percentage points. Another way to interpret the model estimates is that a coefficient of 0.1 means that roughly 10% of the post-treatment intention is explained by the treatment and roughly 90% of the post-treatment intention is explained by the pre-treatment intention, captured by the estimate of β_1.

To maximize efficiency and statistical power, we estimate models pooling all three issues together. To ensure that our results are consistent across issues, we also estimate simple OLS regression models for each issue *j* separately:

$$Y_i = \alpha + \beta_1 P_i + \beta_2 X_i + \varepsilon_i \tag{2}$$

Results

Descriptive Results

Before presenting any model-based estimates, we visually present basic descriptive results which illustrate the core of our main results. The goal here is to present the treatment effects in a straightforward manner without an immediate concern for statistical significance (which we show below). The dependent variable of interest is the change in policy support between the post-treatment and pre-treatment measure. In Figure 1, we plot the mean change in individual-level policy support (averaged across the three issues) for each of the thirteen treatment values (20-80 in increments of five) for both treatment types: polls and probabilities. We also plot the linear relationship between treatment value and individual-level policy support change. Figure 1a illustrates that as the percentage of people supporting the policies in the poll treatment increases from 20% to 80%, individual-level support for the policies increases as well. The linear prediction indicates that showing low levels of public support for the policy (~20%) reduces individual-level support by 5% while polls showing high levels of support have no effect.

This suggests that the bandwagon effect is not driven by the attractiveness of public support but rather the unattractiveness of public opposition. An alternative explanation is that people's priors were that the policies were generally popular, which would mean that information showing them to be popular would have little additional effect on support while information showing them to be unpopular would suppress support.[14] However, in the follow-up study we asked people to report the percentage of the general population they *thought* supported each of three issues. Whereas people thought removing troops from Afghanistan was highly popular (mean perceived support: 70.1%), this was not the case for the free trade (49.4%) and public financing issues (39.2%). The average perceived support across issues was 52.9%, right in the middle of the range of possible treatment values.[15] Hence, polling information can reduce support for policies but cannot increase support.

We find a relatively weak relationship between the treatment value and the dependent variable for the probability treatment (see Figure 1b). The linear fit between the treatment value bins and change in individual-level support is basically flat. Hence, while polls can generate bandwagon effects, probabilities do not. We discuss potential reasons why this may be the case in the conclusion.

Model-Based Estimates

We first present estimates of model (1) for each type of polling information in the "All Issues" columns (columns 1 and 5) of Table 1. For the poll treatment, the relationship between the treatment value and individual-level policy support is positive and statistically significant

[14] We did not measure prior perceptions of public support for the policies because we were concerned that this would prime respondents and blunt the effects of the treatment information.

[15] See Online Appendix 2 for more information on people's prior expectations of public support for the three policies. Interestingly, if we regress individual-level support for the policies against prior beliefs of public support, we obtain statistically significant and substantively large coefficient estimates for all three issues (Afghanistan: β = 0.74, p <.001; free trade: β = 0.93, p <.001; Afghanistan: β = 1.00, p <.001). Of course, these observational results confound the bandwagon effect with the false consensus effect, underscoring the need for an experiment that randomly assigns public support levels.

(see left hand side of Table 1). The coefficient estimate of β_2 is 0.135 (p = 0.001, two-tailed), indicating that moving from 20% to 80% general public support for a policy is associated with an 8.1% increase in individual-level support (60 x 0.135), or about 8.1 points on the 100-point scale. This is substantively significant given that the standard deviation of policy support across all treatment types is approximately 33 points. Conversely, the probability treatment had a substantively small and statistically insignificant impact (β_2 = 0.011, p = 0.78) on policy support (see right hand side of Table 1).[16] Moving across the full range of the treatment value only increased individual-level support for the policies by 0.7 percentage points.[17] We obtained similar results from fixed effects models: poll treatment (β_2 = 0.142, p = 0.001) and probability treatment (β_2 = 0.017, p = 0.71). Columns (2)-(4) and (6)-(8) of Table 1 present estimates from model (2) for each of the three issues separately. While the effect sizes for the probability treatment are inconsistent, the effect sizes for the poll treatment are consistently positive.[18]

Conclusion

This paper shows that polls, by directly influencing individual-level support for policies, can be self-fulfilling prophecies and produce opinion cascades. That conformity pressures can suppress minority opinion may seem disheartening to normative conceptions of democracy (Noelle-Neumann 1974). On the other hand, we found no evidence that information about probabilities, which have become increasingly popular and ubiquitous in recent years, influences

[16] A model including an interaction term between treatment value and treatment type shows that the treatment effect for the poll treatment was significantly greater than the treatment effect for the probability treatment (p = .024).
[17] We also assessed whether there was non-linearity in the treatment effects by estimating nested general additive models (GAMs) (Hastie and Tibshirani 1990) and found that including non-linear terms did not improve model fit beyond a linear specification.
[18] The effect sizes appear to be a bit larger for the free trade issue. To assess whether the effect size for free trade was significantly different than the effect sizes for the other two issues, we estimated regression models for each treatment type including dummy variables for issues and interaction terms between these dummies and the treatment value. None of the three linear combinations testing differences across effect sizes between issues was significant for the probability treatment. Only one of the three differences between issues was significant for the poll treatment (free trade vs. public financing).

people's attitudes in a similar fashion.

The fact that the polling information was impactful while the probability treatments were not suggests that the mechanisms of normative and informational social influence are more powerful than cognitive dissonance reduction in this domain. Polls more directly provide people with information about what their fellow citizens think, meaning that it is easier for people to pick up social cues from such data. On the other hand, probabilities more accurately tell people what is going to happen and therefore potentially activate processes to reduce cognitive dissonance. Yet, we find that people do not change their opinions to be in line with an expected reality.

Our design can be extended to address other substantive questions as well. Although we did include some filler questions, it remains an open question whether the effects of the treatment would sustain across a longer time period. We also only treated people once with the treatment information while in the real world people are exposed multiple times. Additionally, whereas we only presented one poll without disconfirmatory information, future research can explore how competing (and contradictory) sources influence opinion formation. Finally, our dependent variable of interest in this study was policy attitudes; it would be interesting to see if the power of informational social influence extends to candidate choice in the context of elections.[19]

[19] Of course, one thing that makes this difficult to test in an experimental setting is that people are often exposed to polling information in political campaigns, thereby introducing pre-treatment bias. However, one could potentially explore low information contests where pre-treatment information is not ubiquitous. Additionally, cues such as party identification likely swamp majority opinion in candidate choice. Consequently, primary elections are perhaps the most fruitful domain of study, where party is not an obvious cue voters can use to make decisions.

References

Ansolabehere, Stephen, and Shanto Iyengar. 1994. "Of Horseshoes and Horse Races: Experimental Studies of the Impact of Poll Results on Electoral Behavior." *Political Communication*. 11:413-430.

Ansolabehere, Stephen, and Brian F. Schaffner. 2011. "Does Survey Mode Still Matter? Findings from a 2010 Multi-Mode Comparison." SSRN Working Paper 1868229.

Aronson, Elliot, Timothy D. Wilson, and Robin M. Akert. 2005. *Social Psychology*. 5th edition. Upper Saddle River, NJ: Prentice Hall.

Asch, Solomon E. 1951. "Effects of Group Pressure upon the Modification and Distortion of Judgment." In *Groups, Leadership, and Men*. Ed. H. Guetzkow. Pittsburgh, PA: Carnegie Press.

Bartels, Larry. 1988. *Presidential Primaries and the Dynamics of Public Choice*. Princeton, NJ: Princeton University Press.

Brady, Henry, and Richard Johnston. 1987. "What's the Primary Message: Horse Race or Issue Journalism?" In *Media and Momentum: The New Hampshire Primary and Nomination Politics*. Eds. Gary R. Orren and Nelson W. Polsby. London: Chatham House.

Broh, Anthony C. 1980. "Horse-Race Journalism: Reporting the Polls in the 1976 Presidential Election." *Public Opinion Quarterly*. 44:514-529.

Byers, Dylan. 2013. "Nate Silver: I Might Stop Blogging If It Influences Election Results." *Politico*. 14 Feb 2013. Accessed at http://www.politico.com/blogs/media/2013/02/nate-silver-id-stop-blogging-if-it-influenced-election-157049.html <March 12, 2013>

Callander, Steven. 2007. "Bandwagons and Momentum in Sequential Voting." *Review of Economic Studies*. 74:653-684.

Ceci, Stephen K., and Edward L. Kain. 1982. “Jumping on the Bandwagon with the Underdog: The Impact of Attitude Polls on Polling Behavior.” *Public Opinion Quarterly*. 46:228-42.

Deutsch, Morton, and Harold Gerard. 1955. “A Study of Normative and Informational Social Influences upon Individual Judgment.” *Journal of Abnormal and Social Psychology*. 51: 629-36.

Dizney, Henry F., and Ronald W. Roskens. 1962. “An Investigation of the ‘Bandwagon Effect’ in a College Straw Election.” *Journal of Educational Sociology*. 36:108-114.

Easley, Jonathan. 2012. “GOP Takes Aim at ‘Skewed’ Polls.” *The Hill*. 25 Sep 2012. Accessed from http://thehill.com/homenews/campaign/251413-gop-takes-aim-at-skewed-polls <March 12, 2013>

Goidel, Robert K., and Todd G. Shields. 1994. “The Vanishing Marginals, the Bandwagon, and the Mass Media.” *Journal of Politics*. 56:802-810.

Granberg, Donald, and Edward Brent. 1983. "When Prophecy Bends: The Preference-Expectation Link in U.S. Presidential Elections." *Journal of Personality and Social Psychology*. 45:477-491.

Hastie, Trevor, and Robert Tibshirani. 1990. *Generalized Additive Models*. London: Chapman and Hall.

Kay, Aaron, Maria Jimenez, and John Jost. 2002. “Sour Grapes, Sweet Lemons, and the Anticipatory Rationalization of the Status Quo.” *Personality and Social Psychology Bulletin*. 28: 1300-1312.

Knight, Brian, and Nathan Schiff. 2010. “Momentum and Social Learning in Presidential Primaries.” *Journal of Political Economy*. 118:1110-1150.

Krizan, Zlatan, Jeffrey Mille, and Omesh Johar. 2010. “Wishful Thinking in the 2008 U.S.

Presidential Election." *Psychological Science*. 21:140-146.

Lang, Kurt, and Gladys Engel Lang. 1984. "The Impact of Polls on Public Opinion." *Annals of the American Academy of Political and Social Science*. 472:129-142.

Langer, Gary. 2013. "Poll Tracks Dramatic Rise in Support for Gay Marriage." Accessed at http://www.langerresearch.com/uploads/1147a5GayMarriage.pdf <March 26, 2013>

Lavrakas, Paul J., Jack K. Holley, and Peter V. Miller. 1991. "Public Reactions to Polling News During the 1988 Presidential Election Campaign." In *Polling and Presidential Election Coverage*. Eds. Paul J. Lavrakas and Jack K. Holley. Newbury Park, CA: Sage.

Manski, Charles F. 2004. "Interpreting the Predictions of Prediction Markets." NBER Working Paper No. 10359.

Marsh, Catherine. 1984. "Back on the Bandwagon: The Effects of Opinion Polls on Public Opinion." *British Journal of Political Science*. 15:51-74.

McAllister, Ian, and Donley T. Studlar. 1991. "Bandwagon, Underdog, or Projection? Opinion Polls and Electoral Choice in Britain, 1979-1987." *Journal of Politics*. 53:720-740.

Mehrabian, Albert. 1998. "Effects of Poll Reports on Voter Preferences." *Journal of Applied Social Psychology*. 28:2119-2130.

Morton, Rebecca B., and Kenneth C. Williams. 2001. *Learning by Voting: Sequential Choices in Presidential Primaries and Other Elections*. Ann Arbor, MI: University of Michigan Press.

Morwitz, Vicki G., and Carol Pluzinski. 1996. "Do Polls Reflect Opinions or Do Opinions Reflect Polls?" *Journal of Consumer Research*. 23:53-65.

Mutz, Diana C. 1995. "Effects of Horse-Race Coverage on Campaign Coffers: Strategic Contributing in Presidential Primaries." *Journal of Politics*. 57:1015-1042.

Mutz, Diana C. 1997. "Mechanisms of Momentum: Does Thinking Take It So?" *Journal of*

Politics. 59:104-125.

Mutz, Diana C. 1999. *Impersonal Influence: How Perceptions of Mass Collectives Affect Political Attitudes*. New York: Cambridge University Press.

Navazio, Robert. 1977. "An Experimental Approach to Bandwagon Research." *Public Opinion Quarterly*. 41:217-225.

Noelle-Neumann, Elisabeth. 1974. "The Spiral of Silence: A Theory of Public Opinion." *Journal of Communication*. 24: 43-51.

Pew Forum on Religion and Public Life. 2012. "Changing Attitudes on Gay Marriage." Accessed at http://features.pewforum.org/same-sex-marriage-attitudes/ <March 12, 2013>

Pew Research Center for the People and the Press. 2012. "Continued Majority Support for Death Penalty." Accessed at http://www.people-press.org/2012/01/06/continued-majority-support-for-death-penalty/ <March 12, 2013>

Rickershauser, Jill and John H. Aldrich. 2007. "'It's the Electability, Stupid' – or Maybe Not? Electability, Substance, and Strategic Voting in Presidential Primaries." *Electoral Studies*. 26: 371-380.

Rivers, Douglas. n.d. "Sample Matching: Representative Sampling from Internet Panels." Available at: http://www.polimetrix.com/documents/YGPolimetrixSampleMatching.pdf.

Ross, Lee. 1977. "The False Consensus Effect: An Egocentric Bias in Social Perception and Attribution Processes." *Journal of Experimental Social Psychology*. 13: 279–301.

Sherif, Muzafer. 1936. *The Psychology of Social Norms*. New York: Harper.

Simon, Herbert A. 1954. "Bandwagon and Underdog Effects and the Possibility of Election Predictions." *Public Opinion Quarterly*. 18: 245-253

Sinclair, Betsy and Charles Plott. 2012. "Informing the Uninformed: Do Voters Learn From Pre-

Election Polls?" *Electoral Studies*. 31:83-95.

Vavreck, Lynn, and Douglas. Rivers. 2008. "The 2006 Cooperative Congressional Election Study." *Journal of Elections, Public Opinion & Parties*. 18:355-366.

Wolfers, Justin, and Eric Zitzewitz. 2007. "Interpreting Prediction Markets Prices as Probabilities." NBER Working Paper No. 12200.

Table 1: The Effect of Polling and Probability Information on Individual-Level Policy Support

	(1) All Issues	(2) Afghanistan	(3) Free Trade	(4) Public Financing	(5) All Issues	(6) Afghanistan	(7) Free Trade	(8) Public Financing
	Poll Treatment				*Probability Treatment*			
Treatment Value: β_2	0.135*	0.105	0.225*	0.059	0.011	-0.045	0.107	-0.031
	(0.038)	(0.069)	(0.072)	(0.072)	(0.040)	(0.065)	(0.072)	(0.074)
Pre-Treatment Support: β_1	0.792*	0.842*	0.780*	0.790*	0.774*	0.782*	0.809*	0.791*
	(0.026)	(0.042)	(0.048)	(0.039)	(0.029)	(0.040)	(0.049)	(0.038)
Afghanistan Issue: α_1	7.09*	———	———	———	9.92*	———	———	———
	(1.73)				(1.88)			
Public Finance Issue: α_2	1.20	———	———	———	-0.16	———	———	———
	(1.82)				(1.79)			
Constant	-0.989	4.179	-4.816	4.030	8.363*	20.40*	1.903	9.422*
	(2.988)	(4.246)	(4.484)	(4.090)	(2.794)	(4.362)	(4.759)	(4.275)
R^2	0.652	0.644	0.550	0.643	0.660	0.608	0.534	0.643
N	702	234	234	234	735	245	245	245

******p < .05; +p < .10 (two-tailed)

Notes: Random effects regression coefficients in columns (1) and (5) from model (1). OLS regression coefficients in columns (2)-(4) and (6)-(8) from model (2). Standard errors in parentheses. Standard errors clustered by respondent in columns (1) and (5).

Figure 1: The Effect of Polling and Probability Information on Individual-Level Policy Support (Descriptive Results)

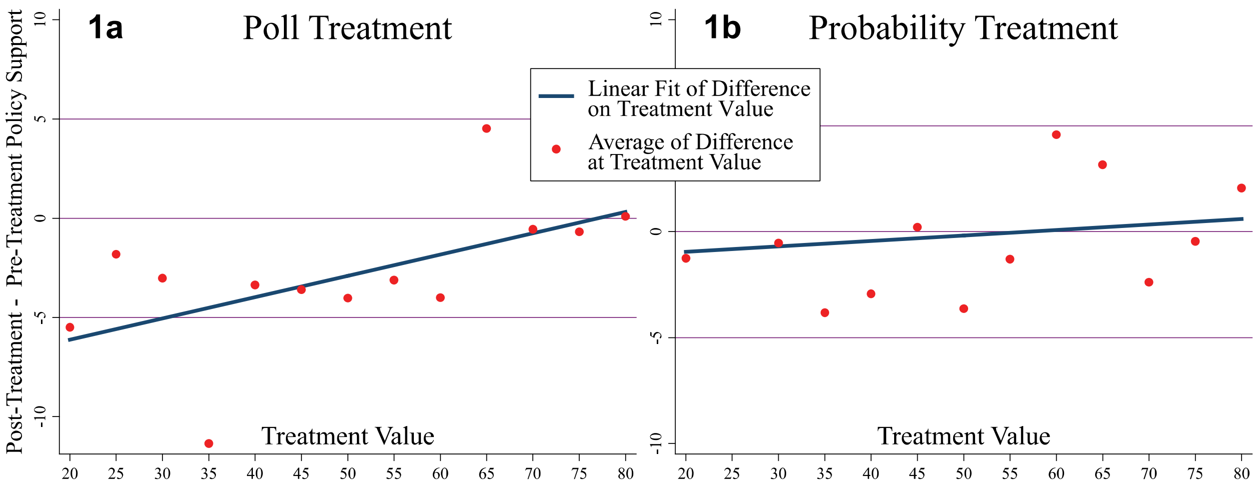

Notes: Points represent bins for each of the 13 possible treatment values ranging from 20% to 80%. Figures plot mean difference in post-treatment and pre-treatment individual-level policy support against level of treatment value.

Appendix 1: Randomization Checks

Appendix Table 1a shows that respondents assigned to each of the four treatment types were statistically and substantively similar across various demographic characteristics. Appendix Table 1b shows that there was no relationship between the treatment value and any pre-treatment demographic characteristic, determined by estimating regressions predicting treatment value with demographic characteristics. The F-statistics of the regressions (i.e., the joint null hypothesis tests that all the coefficients are equal to zero) are all statistically insignificant.

Appendix Table 1a: Randomization Checks for Assignment to Treatment Type

	Treatment Type	
	Poll	Probability
Gender:		
Male	50.4%	48.6%
Female	49.6	51.4
$\chi^2(1) = 0.2$, $p = .69$		
Race:		
White	84.6%	75.5%
Black	5.1	9.8
Hispanic	3.9	9.4
Asian	2.6	2.0
Native American	0.9	0.8
Mixed	1.3	0.8
Other	1.7	1.6
$\chi^2(6) = 10.6$, $p = .10$		
Education:		
Less than HS	3.8%	4.1%
High School	34.2	31.0
Some College	18.4	20.8
2-Year College	8.1	7.4
4-Year College	26.9	26.9
Graduate	8.6	9.8
$\chi^2(5) = 1.0$, $p = .96$		
Age:		
18-29	14.5%	12.7%
30-39	7.3	7.4
40-49	17.1	15.9
50-64	41.5	38.8
65+	19.7	25.3
$\chi^2(4) = 2.3$, $p = .68$		
Party ID:		
Democrat	29.1%	35.9%
Republican	32.5	24.1
Independent	30.3	31.0
Other	1.3	2.0
Not Sure	6.8	6.9
$\chi^2(12) = 5.2$, $p = .27$		

Appendix Table 1b: Randomization Checks for Assignment to Treatment Support Value

	Policy Issue		
	Afghanistan	Free Trade	Public Financing
Gender	-2.313	0.714	1.437
	(1.744)	(1.751)	(1.735)
Race	-1.327	-2.441	-2.997
	(2.199)	(2.207)	(2.188)
Education	-2.570	-2.826	-1.579
	(2.907)	(2.918)	(2.892)
Age	0.565	-0.599	-0.563
	(0.684)	(0.686)	(0.680)
Democrat	1.697	-2.689	-1.932
	(2.056)	(2.064)	(2.045)
Republican	-2.994	-1.029	3.236
	(2.157)	(2.165)	(2.146)
Constant	50.73^{***}	55.75^{***}	52.91^{***}
	(3.527)	(3.540)	(3.508)
R^2	0.015	0.011	0.017
F	1.17	0.86	1.35
p-value	0.32	0.53	0.23
N	479	479	479

**p* < .05; +*p* < .10 (two-tailed)

Note: OLS regression coefficients with standard errors in parentheses.

Appendix 2: Full Question Wordings and Treatment Information

Dependent Variable Questions

Suppose that there was a national referendum on American policy in Afghanistan and you were in the voting booth casting a ballot on the referendum. If you were voting directly on whether or not the U.S. should meaningfully reduce the number of troops in Afghanistan by June 30, 2012, what is the probability that you would vote to reduce the number of troops in Afghanistan?

Suppose that there was a national referendum on American trade policy and you were in the voting booth casting a ballot on the referendum. If you were voting directly on whether or not the U.S. should sign more free trade agreements with North, Central, and South American countries, what is the probability that you would vote for more free trade agreements?

Public financing of state elections is when the government pays for the cost of campaigning for various state offices, rather than the campaigns relying on donations from the general public, corporations, or unions. Suppose that there was a state referendum on campaign finance policy and you were in the voting booth casting a ballot on the referendum. If you were voting directly on whether or not to publically finance elections in your state, what is the probability that you would vote for the public financing of elections?

Poll Treatments

Below is the percentage of Americans who support a meaningful reduction in U.S. troops in Afghanistan by June 30, 2012. *This value is created by aggregating the best available polls.*

Below is the percentage of Americans who support more free trade agreements with North, Central, and South American countries. *This value is created by aggregating the best available polls.*

Below is the percentage of Americans who support public financing of elections. *This value is created by aggregating the best available polls.*

Probability Treatments

Below is the likelihood of there being a meaningful reduction of U.S. troops from Afghanistan by June 30, 2012. The U.S. currently has 100,000 troops in Afghanistan and a meaningful reduction is defined as 80,000 or less troops left. *This value is created by aggregating the best available forecasts.*

Below is the likelihood that public financing will be in place in 10 or more U.S. states by January 1, 2016. There are currently 6 states with some form of public financing. *This value is created by aggregating the best available forecasts.*

Below is the likelihood of the United States having 15 or more free trade agreements with North, Central, and South American countries by January 1, 2020. The U.S. currently has free trade agreements with 10 of the 34 other countries in North, Central, and South America. *This value is created by aggregating the best available forecasts.*

Appendix 3: Graphical Presentations

Pre-Treatment Dependent Variable

Suppose that there was a national referendum on American policy in Afghanistan and you were in the voting booth casting a ballot on the referendum.

If you were voting directly on whether or not the U.S. should meaningfully reduce the number of troops in Afghanistan by June 30, 2012, what is the probability that you would vote to reduce the number of troops in Afghanistan?

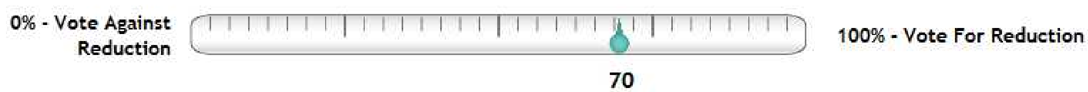

Graphical Presentation for Post-Treatment Dependent Variable

Suppose that there was a national referendum on American policy in Afghanistan and you were in the voting booth casting a ballot on the referendum.

If you were voting directly on whether or not the U.S. should meaningfully reduce the number of troops in Afghanistan by June 30, 2012, what is the probability that you would vote to reduce the number of troops in Afghanistan?

Click on the thermometer to give your rating

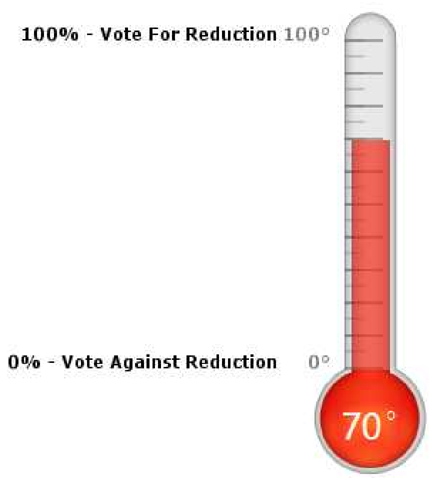

Treatment

Below is the likelihood that public financing will be in place in 10 or more U.S. states by January 1, 2016. There are currently 6 states with some form of public financing. *This value is created by aggregating the best available forecasts.*

25%

Appendix 4: Distributions of Pre- and Post-Treatment Policy Support

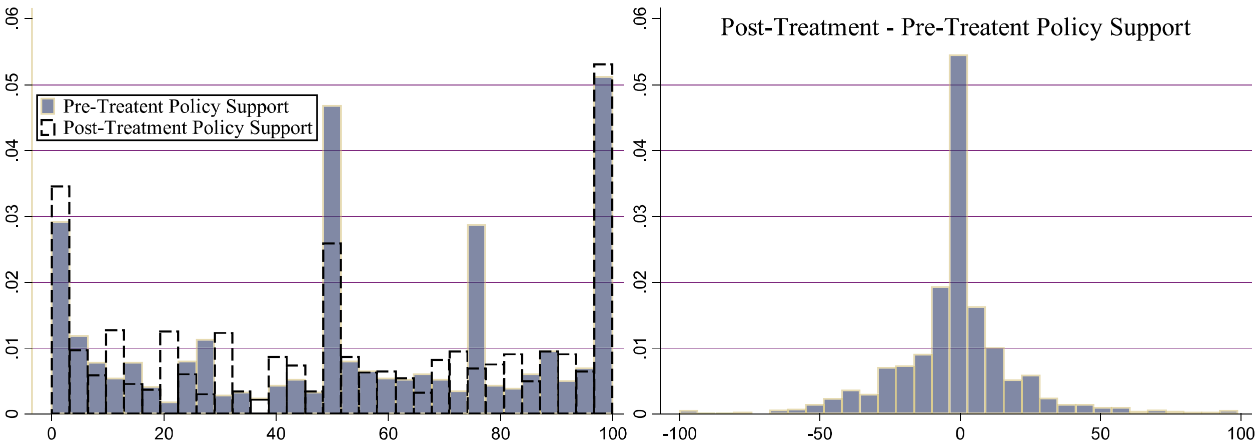

Online Appendix

for

Are Polls and Probabilities Self-Fulfilling Prophecies?

Online Appendix 1: Mean Values of Continuous Vote Intention Dependent Variable by Rating Scale Categories

	Afghanistan	Free Trade	Public Financing
Strongly support	94.2	91.9	95.9
Somewhat support	76.2	66.8	64.0
Neither support nor oppose	55.5	46.0	46.2
Somewhat oppose	36.1	25.6	21.9
Strongly oppose	21.7	4.0	5.8
N	500	496	499

Note: Continuous vote intention variable measured on 0-100 scale. Question wordings for Likert scales are: (1) “How much do you support or oppose withdrawing all troops from Afghanistan by June 30, 2013?”; (2) “How much do you support or oppose the U.S. signing more free trade agreements with North, Central, and South American countries?”; (3) “How much do you support or oppose public financing for state elections?” The question wordings for the continuous vote intention dependent variables are the same as in Appendix 1 in the manuscript.

Online Appendix 2: Prior Expectations of Public Support

	Mean	Std. Dev.	Median	25th Pcntile	75th Pcntile	N
Afghanistan	70.1	16.0	71.0	60.0	80.0	500
Free Trade	49.4	17.1	50.0	40.0	60.0	499
Public Financing	39.2	24.0	35.0	20.0	56.0	499
Average Across Issues	52.9	12.1	53.0	43.7	61.0	498

Note: Prior expectations variable measured on 0-100 scale. Question wordings are: (1) "What percentage of Americans do you think support withdrawing all troops from Afghanistan by June 30, 2013?"; (2) "What percentage of Americans do you think support the U.S. signing more free trade agreements with North, Central, and South American countries?"; (3) "Public financing of state elections is when the government pays for the cost of campaigning for various state offices, rather than the campaigns relying on donations from the general public, corporations, or unions. How much do you support or oppose public financing for state elections?"