The Subtle Psychology of Voter Turnout

Daniel G. Goldstein,[1][*] Kosuke Imai,[2] Anja S. Göritz[3]

[1]London Business School, UK
[2]Princeton University, USA
[3]University of Erlangen-Nuremberg, Germany

[*]To whom correspondence should be addressed: dgoldstein@london.edu

March 26, 2007

Randomized experiments conducted during the 2006 US midterm election and the 2005 German federal election were analyzed in order to estimate the turnout effects of two simple treatments: asking people *if* and *how* they intend to vote.

Since World War II, in over 1,600 national elections in 170 independent states, voter turnout rates have averaged about 65% of the voting age population (*1, 2*). Policy makers in 18% of these democracies have deemed electoral participation important enough to justify compulsory voting laws, under which non-voters can face fines and other punishments. Recently, the US Congress authorized 3.9 billion dollars for the Help America Vote Act, and state governments have invested in expanding early-voting methods, which accounted for roughly 20% of the votes cast in the 2004 election (*3, 4*). Worldwide, rewards for voters have included tax breaks, job opportunities, scholarships, and even high-stakes lotteries (*1*). Alongside government campaigns are partisan ones: US Democrats and Republicans spent roughly 100 million dollars on selective turnout programs in the 2000 election alone (*5*).

What drives voter turnout? Political theory speaks of the costs and benefits of voting and the slight probability that one's vote will be decisive (*6*). In practice, these and other variables appear in policies that target two causes of weak participation: low motivation and high obstacles (*7*). Motivation-focused initiatives aim to impart the desire to vote by invoking the importance or closeness of an election, a voter's sense of duty, rewards, punishments, or social comparisons. Obstacle-focused policies aim to make voting easier, such as by introducing same-day or automatic registration, voting by mail, or early in-person voting.

If voting is largely influenced by motivations and obstacles, policy makers might take inspiration from psychological research on goal attainment, which has revealed the strong effects of two simple treatments. In this paper, we demonstrate that simply asking people *if* and *how* they intend to vote can increase turnout.

The technique of asking people if they intend to vote comes from research on attitude accessibility and self-fulfilling prediction. In what is called the *mere measurement* or *question-behavior* effect (*8,9*), people become more likely to perform certain actions if they are first asked whether they expect to perform them. That is, merely measuring intentions changes behavior. One surprising study found that asking people whether they intended to buy an automobile increased their chances of doing so (*9*).

Why does mere measurement work? One important literature suggests that people who make forecasts about the future may alter their behavior to make the predictions come true (*10*). An emerging and complementary view is that when people answer questions about intentions, their underlying attitudes become concrete and readily accessible (*9, 11*). For this reason, questions can be polarizing. If underlying attitudes are positive, mere measurement can increase the likelihood of performing an action – if negative, the same questions can decrease it. If attitudes toward electoral participation are generally positive, assessing intentions may turn voting into a goal.

2

Eighty years of research has looked at the effect of polls, questions, and surveys on voter turnout ((*12*) for a review, see (*13*)). For instance, twenty years ago, Greenwald and colleagues (*14*) found some promising but tentative results on a sample of sixty college students. However, theirs, like many if not all others in the literature, was not a pure test of mere measurement, as participants were not just asked about intentions, but were also provided information about where and how to register to vote if they did not know. Over the last 80 years, research has left an unclear picture of the effects of mere-measurement on voting due to mixed results and some methodological controversies (*13, 15*). Part of the variation in results may be due to the variety of populations, instruments, and historical periods studied. Additional variation may be due to the way experiments have mixed mere-measurement treatments with related political questions and even practical information for voters.

The second technique, asking people how they intend to vote, comes from research on *implementation intentions*, which are simple plans that help people overcome obstacles en route to goal attainment. The effects of implementation intentions have been estimated in over 100 policy-relevant studies on exercising, recycling, smoking, and beyond, however, the link to voting has not been investigated in the literature (*16*).

How do implementation intentions work? These plans are hypothesized to lead one to direct resources (such as time and attention) toward a target goal, and away from competing goals when they inevitably arise. Furthermore, implementation intentions might make one aware of goal-realization opportunities that would otherwise go unnoticed (e.g., noticing registration offices near work), and help automate responses to foreseeable obstacles (e.g., identifying a means of backup transportation to the polls) (*16*).

We illustrate the application of mere measurement and implementation intentions through randomized experiments, analyzed in order to estimate causal effects on voter turnout in two national elections: the 2006 US Midterm Election and the 2005 German Federal Election.

3

We present here a summary – methodological details are found in the appendix. In the US study, $1,968$ voting-age members of a nationwide research panel were invited to take part in a brief survey approximately two months before the election. In it, a mere measurement group was asked about intentions to vote, and an implementation intentions group was additionally asked to formulate simple plans to vote. A control group completed a filler task to equalize time. The crucial difference with the German study, which invited $1,426$ people, is that it took place 1 to 4 days before the election, presumably leaving treatments fresh in the minds of participants. After the election, a follow-up study measured turnout (election day voting and early voting). As detailed in the appendix, we use a diary technique to validate US election day voting and make statistical adjustments for non-random dropout in both experiments.

The experiments pose novel theoretical and applied questions. Will implementation intentions have an effect above that of mere measurement? Will the two treatments be effective on *one-shot goals* (*17*) that can be realized only on one day (e.g., voting on Election Day) and *open-ended goals* that can be realized on many possible days (e.g., early and postal voting)? For both types of goals, do mere measurement and implementation intentions treatments fade over periods of days or months?

The methodological challenge of these experiments is that not all people participate in follow-up studies. Although the problem of such drop-out is common in survey experiments (*18*), the main concern here is that the drop-out probability may directly depend on the voting behavior: those who vote may be more likely to report their voting behavior (*19*). The methodology we develop in the appendix material allows for the drop-out probability to directly depend on the (possibly unobserved) outcome variable, as well as various pre-treatment covariates measured for each subject, thereby yielding valid statistical inference (*20*).

[Figure 1 about here.]

4

The results of the experiments are displayed in Figure 1. For the *open-ended* goal of early (e.g., postal) voting, mere measurement treatments given two months in advance (US study) had moderate positive effects on turnout, a finding consistent with studies showing that mere measurement treatments can impact the probability of undertaking an action (such as purchasing a computer) on any day within a window of several months (*9*). For early voting, estimated implementation intentions effects were similar to those of mere measurement, though 2.7 percentage points greater.

For the *one-shot* goal of election-day voting, mere measurement was only effective when it was administered days (Germany), but not months (US) in advance. Implementation intentions treatments, in contrast, held their effectiveness for both near and distant races.

Policy makers seeking to craft inexpensive and unobtrusive ways of increasing turnout will note that mere measurement was about as effective as implementation intentions for all conditions except distant, one-shot goals. However, 14% of those who were assigned to the US implementation intentions group refused to write plans while only 3% of the mere measurement group declined to provide intentions. If the compliance rate was higher in the implementation intentions group, the actual treatment effect might have been larger. Nevertheless, our estimates are relevant to policy since a similar pattern of noncompliance is expected to occur in practice.

Our study contributes to a growing body of research demonstrating that policies can benefit from working in concert with psychological mechanisms. People's preference for default options, for instance, can lead to increased membership in organ donor pools (*21*), and participation in retirement savings plans (*22*). While some policies benefit from a tendency toward inaction, others must help people to act. To construct effective campaigns and messages, policy makers might consider addressing voting as a goal, one that is aided by stating intentions and making plans.

Appendix

2006 US Midterm Election Experiment

Participants were 2,469 members of the London Business School Online Laboratory Panel, selected from the database for being US citizens of voting age. The experiment consisted of two phases. The first phase ran from September 15th to 25th, 2006, some seven to eight weeks before the election of November 7th, 2006. The second phase took place from November 8th to 11th, 2006, one to four days after the election. Before the first phase, participants were randomly assigned to three groups of equal size; a control group, mere measurement (MM) group, and implementation intentions (II) group. Members of all groups were sent identical emails inviting them to participate in an online study on "decision making." Every participant was offered payment of one US dollar in addition to entry in a lottery in which one participant would receive 100 dollars, five participants would receive 20 dollars, and 10 participants would receive 10 dollars.

To improve the efficiency of the resulting causal estimates, we used a matched-pair design where complete randomization of the treatments was conducted within each group of three observations with similar characteristics (*18, 23*). These randomized groups are formed based on the Mahalanobis distance of the selected pre-treatment variables, which is a common metric used to measure the degree of similarity among observations in terms of those multi-dimensional characteristics. The variables we used are years in residence, gender, age, age squared, marital status, employment status (full-time or not), whether a respondent has at least one child, whether a respondent is most fluent in English, annual income, and years of education.

From this initial sample, we first exclude 62 invitees who had non-functioning email addresses or inconsistent demographic information. The group sizes for the control, MM, and

II groups are reduced to 800, 809, 798, respectively. (As a convention, groups will always be listed in this order.) Next, we exclude those who did not begin the experiment at all, measured by agreeing to a consent page that is identical across conditions. This yields group sizes of 400, 430, and 379 for the final sample we analyze. The above exclusion criteria are based on the pre-treatment information that could not be affected by the treatment conditions, thus any variation across the groups is assumed to have arisen at random. As expected, the observed covariates are well balanced between the three groups. For example, Pearson's χ^2 tests show that the observed differences in gender, marital status, and employment status are not statistically significant.

During the first phase, the control group provided basic demographic information, then completed a five-minute filler task, used to equalize task time with the other conditions. The MM group was identical to the control, except that before the filler task, it was asked to indicate the strength of voting intentions by rating the following statements from Dholakia and Bagozzi (*17*) on a 9 point scale ranging from "I agree completely" to "I disagree completely": "I intend to vote in the upcoming US Midterm Election," "I am very committed to voting in the upcoming US Midterm Election," and "It would not take much for me to abandon my goal of voting in the upcoming US Midterm Election."

The II group was the same as the MM group but without the filler task and with the addition of items used to elicit implementation intentions, also adapted from Dholakia and Bagozzi. The items were the following three questions, each followed by a text entry box on the Web form: 1) "If you are not registered to vote, please write a few sentences answering the following questions by listing specific steps. How will you find out about registering? When will you find out about registering? When will you register? Where will you register? How will you register?", 2) "Listing specific steps, please write a few sentences answering the following questions. How will you vote (in person or by mail)? Where will you look for information on voting? If you vote in person, how will you find out where to vote? When will you find out your voting location?

7

If you will vote by mail, how will you find out about postal voting? When will you find out about postal voting?", and 3) "Listing specific steps, please write a few sentences answering the following questions. If you vote in person, what time of day will you go to vote? Where will you vote? How will you get to your voting location? If you vote by mail, when will you mail your ballot? Where will you mail it from?" Each of the three items was followed by a question asking for a contingency plan of the following kind: "Sometimes a good plan prevents a difficult situation from happening. If you are not registered to vote, please write a few sentences about what might go wrong with this plan for registering to vote. What will you do if these obstacles arise? What can you do to prevent them from arising?" People not intending to vote were instructed to answer as if they did intend. At no point in the experiment was the possibility of a follow-up study mentioned.

To avoid the problem of deception in self-reports (*24*), the following diary method was used to measure in-person voting in the four days after the election, some two months after the treatment. Diary methods offer a complementary line of evidence to the method of checking votes against electoral rolls, which has some drawbacks, such as limiting the sample to people already in voter registration databases (not allowing for the possibility of registering as a result of the experiment), limiting the sample to select counties (as opposed to the entire country), and missing "lost votes" (due to faulty equipment, registration mix ups, polling place operations, lost absentee ballots) of which there were 4 to 6 million in the 2000 US Presidential election (*25*).

All the original invitees received a seemingly unrelated boilerplate email invitation about a "short research study on memory". Upon starting the memory study, participants were informed that they would be asked to remember, in as much detail as possible, what they did hour by hour on a particular day. To illustrate the task, they were shown an example diary with items such as "7AM - Drove to the Lazy Daisy Cafe" and "12PM - Had a cheese sandwich for lunch in the cafeteria". Participants were then given a blank diary with fields for each hour from 5AM

8

through 11PM, and were told which day to remember. Each person actually received the same day, Tuesday, and no mention was made of it being the midterm election day. After completing the diaries, they were submitted irrevocably online. At this point, participants were asked if they voted in person on Tuesday, voted before Tuesday (by post or early in-person voting), or did not vote. All participants were paid one dollar and fifty cents for participation, plus entry in a small stakes lottery. In the analysis phase, respondents were coded as having voted in person if and only if they listed having voted in their diaries. The number of participants who agreed to participate in the second phase in the control, MM, and II groups was 242, 260, and 207, respectively. In this experiment, the possibility of non-random dropout is an important methodological issue and is addressed in our statistical analysis.

2005 German Federal Election Experiment

Participants were members of two German Web panels, Promio.net and WiSo-Panel (*26*). The experiment consisted of two phases. The first phase ran from September 13th to 17th, 2005, one to five days before the election of September 18th, 2005. The second phase took place from September 19th to 22nd, 2005, one to four days after the election. Before the first phase, panel database members were randomly assigned to the control group, mere measurement group, or implementation intentions group. Members of all groups were sent identical emails inviting them to participate in an online study on "decision making". For an analysis not relevant to the focus of this paper, participants were either paid by the panels, paid 1.5 Euros by our lab, or decided to take part in the experiment without pay. Web links in 252, 584, and 590 invitations were clicked on. (The latter two treatment conditions were allowed to gather more responses in the interest of detecting differences between them.) Of these, 249, 579, and 586 agreed to a consent page. This is the sample we analyze in this paper.

During the first phase, the control group provided basic demographic information. The MM

group was identical to the control group except that it was given a yes-no item about intention to vote. The II group was identical to the MM group, except that it was given two additional items. The first asked those intending to vote to list one main obstacle that might prevent them from voting. The second requested that they write a plan they could use to overcome this obstacle should it arise. In the second phase, after the election, participants were asked whether they voted in person, voted by mail, or did not vote. All respondents were invited to participate after the election, and 204, 485, and 512 individuals consented. Similar to the US election experiment, the possibility of non-random dropout is addressed in our statistical analysis.

Statistical Methodology

Finally, we describe the statistical methodology used to analyze our two experiments. Let T_i be the multi-valued treatment variable where $T_i \in \mathcal{J} = \{0, 1, \ldots, J-1\}$. In our experiments, J is equal to 3 – the control, MM, and II groups – but we maintain the general notation throughout this section. We develop our methodology under the *potential outcomes* framework of causal inference commonly used in statistics (*27*). Let $Y_i(t) \in \{0, 1\}$ denote the *potential* turnout variable of voter i under the scenario that the voter receives the treatment value $t \in \mathcal{J}$. For example, $Y_i(0) = 1$ means that if voter i is assigned to the control group, he/she would vote in the election. For each voter, we only observe one of the J potential outcome variables, and this observed outcome variable is denoted by $Y_i = \sum_{j=0}^{J-1} \mathbf{1}(T_i = j) Y_i(j)$ where $\mathbf{1}(\cdot)$ is the indicator function. The fundamental problem of causal inference is that any causal effect involves at least two potential outcomes, e.g., $Y_i(1) - Y_i(0)$, and hence can never be directly observed.

To address the issue of non-random dropout in our experiments, we use $R_i(T_i) \in \{0, 1\}$ to represent the *potential* binary recording variable which equals 1 if voter i would report actual voting behavior in the post-election survey after receiving the treatment value of T_i (and 0 otherwise). There are J potential recording variables for each voter, but as with the potential

turnout variable, we only observe one of them. We denote the observed recording variable $R_i = \sum_{j=0}^{J-1} \mathbf{1}(T_i = j) R_i(j)$. We also observe some covariates for each participant such as age and gender, which are measured in the pre-election survey and denoted by X_i.

The quantity of interest we report in Figure 1 is the *sample average treatment effect* or ATE for each treatment,

$$\tau_j \;\; = \;\; \frac{1}{n} \sum_{i=1}^{n} Y_i(j) - Y_i(0), \tag{1}$$

for $j = 1, \ldots, J - 1$. Since the treatment assignment is randomized, the treatment variable T_i is independent of all potential outcomes, the estimation of the ATE is straightforward if the outcome variable is observed for every participant. However, in our experiments, some voters did not participate in the post-election survey and hence the outcome variable Y_i was not recorded for them. This means that the standard complete-case analysis yields a biased estimate of the ATE unless the drop-out behavior is generated *completely at random*, which is unlikely in this case.

To make more credible inferences, one possibility is to assume that the data are *missing at random* (MAR). That is, a voter's decision to report voting behavior the post-election survey may depend on the treatment he/she received, but given the received treatment and observed covariates, the decision is assumed to be independent of whether he/she voted in the election, i.e., $\Pr(R_i(j) = 1 \mid T_i = j, Y_i(j) = 0, X_i) = \Pr(R_i(j) = 1 \mid T_i = j, Y_i(j) = 1, X_i)$ for all $j \in \mathcal{J}$. Although the MAR assumption may be reasonable in some situations, it is problematic in our experiments because those who voted in the election are more likely to report their voting behavior (*19*). Thus, we allow for the possibility that a voter's decision to answer the post-election survey in our experiments may depend on his/her voting behavior itself by assuming that treatment assignments affect the response decision only *indirectly* through the voting behavior. The filler tasks given to the control and MM groups in the pre-election

survey are intended to equalize the survey time, thereby minimizing the possibility of the *direct* effects of the treatments on the dropout mechanism. Formally, the assumption is called the non-ignorability (NI) assumption and can be written as,

$$\Pr(R_i(j) = 1 \mid T_i = j, Y_i(j) = k, X_i) \;\; = \;\; \Pr(R_i(j') = 1 \mid T_i = j', Y_i(j') = k, X_i), \quad (2)$$

for all $j \neq j'$ and $k = 0, 1$. It has been shown that under the NI assumption, the average treatment effects are identified (*20*) (see also (*28*)).

To estimate the ATE, we conduct a Bayesian inference under the NI assumption in Equation 2 by following the estimation strategies described in (*20*). In particular, we model the turnout and missing data mechanism using the two probit regressions,

$$p_j(z) \;\; \equiv \;\; \Pr(Y_i = 1 \mid T_i = j, Z_i = z) = \Phi(\alpha_j + \beta^\top z), \quad (3)$$

$$r_k(x) \;\; \equiv \;\; \Pr(R_i = 1 \mid Y_i = k, X_i = x) = \Phi(\delta_k + \gamma^\top x), \quad (4)$$

for $j \in \mathcal{J}$ and $k = 0, 1$ where Z_i is a subset of X_i as explained below and $\Phi(\cdot)$ represents the cumulative distribution function of the standard normal random variable.

Equation 2 suggests that in order for the NI assumption to hold, it is important to control for relevant confounding covariates in the response model of Equation 4 (*18*). The key variable we use for this purpose is the vote intention from the pre-election survey, which was measured for those in the MM and II groups. We also include an indicator variable for the voters whose vote intention variable was not observed (i.e., those in the MM and II groups who did not answer this question as well as everyone in the control group). These variables are not included in the outcome model of Equation 3, however, because they constitute a part of the treatments of interest. Furthermore, we include several pre-treatment control variables in both the outcome and response models. For the US experiment, we include gender, age, age squared, education, marital status, number of years in residence, employment status, log income, and an indicator variable

about whether any of the pre-treatment covariates is missing. For the German experiment, we include gender, age, age squared, and the missing covariate indicator variable.

Our inference is based on the following complete-data likelihood function,

$$\prod_{i=1}^{n} \left[\prod_{j=0}^{J-1} \left\{ p_j(Z_i)^{Y_i}(1 - p_j(Z_i))^{(1-Y_i)} \right\}^{\mathbf{1}(T_i=j)} \right] \left\{ r_1(X_i)^{R_i}(1 - r_1(X_i))^{(1-R_i)} \right\}^{Y_i}$$

$$\times \quad \left\{ r_0(X_i)^{R_i}(1 - r_0(X_i))^{(1-R_i)} \right\}^{1-Y_i}, \tag{5}$$

from which the observed-data likelihood function can be computed by integrating out the missing data, i.e., Y_i with $R_i = 0$. Finally, we assume conjugate prior distributions with large variances for the coefficients, i.e., $(\alpha_j, \beta, \gamma) \sim \mathcal{N}(0, 100I)$ and $\delta_k \sim \mathcal{N}(0, 10)$ where I represents the identity matrix. A Markov chain Monte Carlo algorithm is constructed to sample from the posterior distributions. The algorithm is based on a standard Gibbs sampler for the probit regression, but we apply marginal data augmentation to improve its convergence (*29*). For each analysis, a total of one million draws are obtained and the inference is based on every 10th draw of the second half of the chain. The standard diagnostics tools indicate that a satisfactory degree of convergence is attained.

References and Notes

1. A. Ellis, M. Gratschew, J. H. Pammett, E. Thiessen, *Engaging the Electorate: Initiatives to Promote Voter Turnout From Around the World* (International Institute for Democracy and Electoral Assistance, Stockholm, 2006).

2. R. L. Pintor, M. Gratschew, *Voter turnout since 1945 : a global report* (International Institute for Democracy and Electoral Assistance, Stockholm, 2002).

3. United States Election Assistance Commission, *A Summary of the 2004 Election Day Survey. How We Voted: People, Ballots, and Polling Places* (United States Election Assistance Commission, New York, 2005).

4. National Annenberg Election Survey, Early voting reaches record levels in 2004, national annenberg election survey shows, *Press release*, Annenberg Public Policy Center of the University of Pennsylvania (March 24, 2005).

5. J. Dao, *New York Times* pp. 7 November, sec.A (2000,).

6. A. Blais, *To Vote or not to Vote: The Merits and Limits of Rational Choice Theory* (University of Pittsburgh Press, Pittsburgh, 2000).

7. P. E. Converse, Non-voting among young adults in the united states, In *Political parties and political behavior*, W. J. Crotty, D. M. Freeman, D. S. Gatlin, eds. (Allyn & Bacon, Boston, 1971).

8. D. E. Sprott, *et al.*, *Social Influence* **1**, 128 (2006).

9. V. G. Morwitz, E. J. Johnson, D. Schmittlein, *Journal of Consumer Research* **20**, 46 (1993).

10. S. J. Sherman, *Journal of Personality and Social Psychology* **39**, 211 (1980).

11. V. G. Morwitz, G. J. Fitzsimons, *Journal of Consumer Psychology* pp. 64–74 (2004).

12. H. F. Gosnell, *Getting-Out-the-Vote: An experiment in the stimulation of voting* (University of Chicago Press, Chicago, 1927).

13. C. B. Mann, *The Annals of the American Academy of Political and Social Science* **601**, 155 (2005).

14. A. F. Greenwald, C. Carnot, R. Beach, B. Young, *Journal of Applied Psychology* **72**, 315 (1987).

15. K. Imai, *American Political Science Review* **99**, 283 (2005).

16. P. M. Gollwitzer, P. Sheeran, *Advances in Experimental Social Psychology* **38**, 69 (2006).

17. U. M. Dholakia, R. P. Bagozzi, *Journal of Applied Social Psychology* **33**, 889 (2003).

18. Y. Horiuchi, K. Imai, N. Taniguchi, *American Journal of Political Science* (2007). Forthcoming.

19. B. C. Burden, *Political Analysis* **8**, 389 (2000).

20. K. Imai, Statistical analysis of randomized experiments with nonignorable missing binary outcomes, *Tech. rep.*, Department of Politics, Princeton University (2006).

21. E. J. Johnson, D. G. Goldstein, *Science* **302**, 1338 (2003).

22. B. C. Madrian, D. F. Shea, *The Quarterly Journal of Economics* **116**, 1149 (2001).

23. R. Greevy, B. Lu, J. H. Silber, P. Rosenbaum, *Biostatistics* **5**, 263 (2004).

24. D. Granberg, S. Holmberg, *American Journal of Political Science* **35**, 448 (1991).

25. R. M. Alvarez, S. Ansolabehere, E. Antonsson, J. Bruck, *Voting: What Is, What Could Be* (CalTech/MIT Voting Technology Project, 2001).

26. A. Göritz, Using online panels in psychological research, In *Oxford Handbook of Internet Psychology*, A. Joinson, K. McKenna, T. Postmes, U. D. Reips, eds. (Oxford University Press, Oxford, in press).

27. P. W. Holland, *Journal of the American Statistical Association* **81**, 945 (1986).

28. K. Hirano, G. Imbens, G. Ridder, D. B. Rubin, *Econometrica* **69**, 1645 (2001).

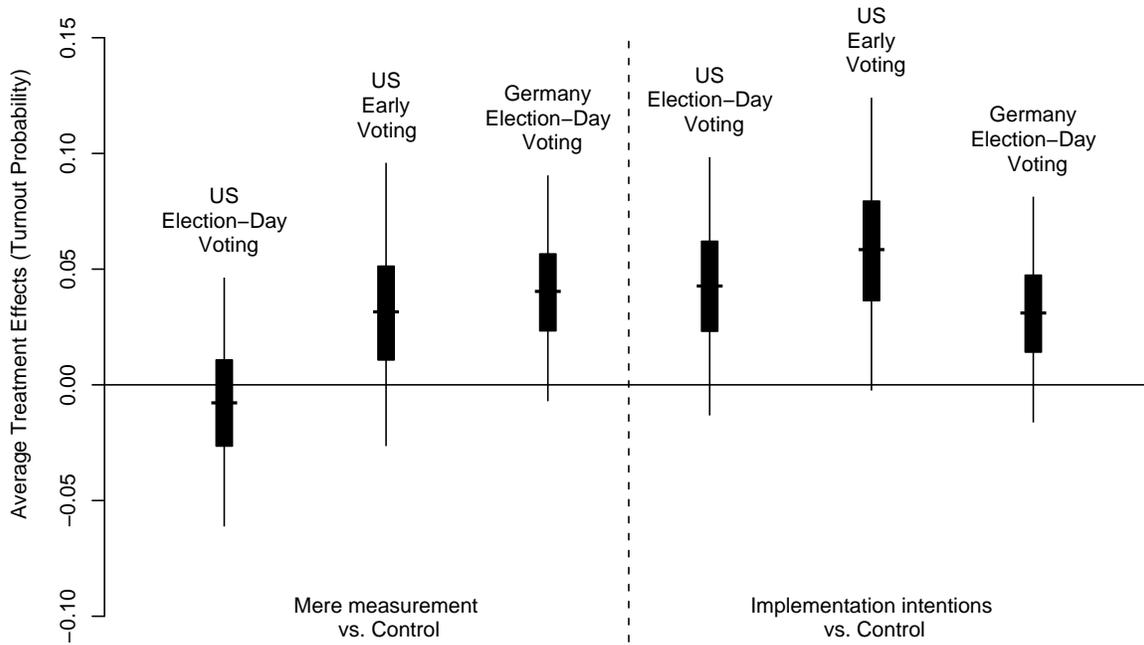29. K. Imai, D. A. van Dyk, *Journal of Econometrics* **124**, 311 (2005).

Figure 1: Estimated Sample Average Treatment Effects of Mere Measurement and Implementation Intentions from the US and German Election Experiments. The figure shows 50% (thick bars) and 95% (thin lines) Bayesian confidence intervals as well as point estimates. For the US experiment, in which treatments came two months before the election, mere measurement does not increase the probability of in-person voting (the estimated average treatment effect or ATE is -0.008 and $\Pr(\text{ATE} > 0) = 0.38$), while it has a modest positive effect on the probability of early voting (ATE is 0.032 with $\Pr(\text{ATE} > 0) = 0.85$). The implementation intentions treatment increases the probability of both in-person voting (ATE is 0.043 and $\Pr(\text{ATE} > 0) = 0.93$) and early voting (ATE is 0.058 and $\Pr(\text{ATE} > 0) = 0.97$). For the German experiment, in which treatments immediately preceded the election, both mere measurement and implementation intentions increase the probability of in-person voting (ATE $= 0.040$ and $\Pr(\text{ATE} > 0) = 0.95$ for mere measurement; ATE $= 0.031$ with $\Pr(\text{ATE} > 0) = 0.89$ for implementation intentions).