

Beliefs and Voting Decisions: A Test of the Pivotal Voter Model

John Duffy* and Margit Tavits†

May 2007

Abstract. We report results from a laboratory experiment testing the basic hypothesis imbedded in various rational voter models that there is a direct correlation between the strength of an individual's belief that his/her vote will be pivotal and the likelihood that individual incurs the cost to vote. This belief is typically unobservable. In one of our experimental treatments we elicit these subjective beliefs using a proper scoring rule that induces truthful revelation of beliefs. This allows us to directly test the pivotal voter model. We find that a higher subjective probability of being pivotal increases the likelihood that an individual votes, but the probability thresholds used by subjects are not as crisp as the theory would predict. There is some evidence that individuals learn over time to adjust their beliefs to be more consistent with the historical frequency of pivotalness. However, many subjects keep substantially overestimating their probability of being pivotal.

* Department of Economics, University of Pittsburgh, Pittsburgh, PA 15260, email: jduffy@pitt.edu.

† Department of Political Science, University of Missouri-Columbia, Columbia, MO 65211, email: tavitsm@missouri.edu.

Why do people vote? While many theories have been offered (for a survey see Dhillon and Peralta 2002), the simplest and most widely used framework is the pivotal voter model (Ledyard 1984; Palfrey and Rosenthal 1983; 1985; see also Downs 1957; Tullock 1967). This model asserts that voters have only instrumental concerns – their motivation is to affect the outcome of the election as opposed to non-instrumental motivations, e.g. warm-glow altruism – and that in making the decision to vote they are rational, self-interested expected payoff maximizers. In particular, people vote if the expected benefit of voting is greater than the cost.

This model has widespread appeal but it is simultaneously the most extensively debated theory in political science (Green and Shapiro 1994, 47-48). The problem is straightforward: the expected benefit calculation involves the use of the voter's probability that she will be pivotal to the election outcome. As in large electorates this probability is very small, rational citizens should not vote. This, however, contradicts the evidence. It is this paradox that feeds the rational choice controversy (Friedman 1995); the paradox argued to have “eaten” rational choice theory (Fiorina 1990, 334). Indeed, the apparent anomaly has led to the search for an extra term – the “D-term” or “a sense of civic duty” – to make voting rational (Riker and Ordeshook 1968). However, this explanation remains theoretically unrewarding (Bendor, Dermeier and Ting 2003).

Given this central controversy in the discipline, it is curious that empirical studies examining the assumptions and predictions of the pivotal voter model are scarce and indirect. Field data can usually provide only weak tests of the model as they pose challenge to measurement and provide little control over extraneous factors (see Levine and Palfrey 2007). Among the difficulties are the unobservability of voters' costs of voting, benefits from an election victory and their beliefs as to whether they will be pivotal to the election outcome – all of which play a critical role in the theory (Green and Shapiro 1994, 47-71).

Undoubtedly, the greatest controversy surrounds the measurement and relevance of the probability of any voter being pivotal – the trademark of the rational choice theory of turnout (Aldrich 1993; Foster 1984; Green and Shapiro 1994, 47-71). Various proxies have been used to measure pivotalness, such as the expected or perceived closeness of the election (Blais et al. 2000; Ferejohn and Fiorina 1975; Foster 1984; see also Matsusaka and Palda 1993 for a review), and the size of the electorate (Hansen, Palfrey and Rosenthal 1987; see also Bendor, Diermeier and Ting 2003, 274-5). However, these proxies have been criticized as being a “far cry” from the actual concept of pivotalness (Aldrich 1993, 259; Cyr 1975, 25; Green and Shapiro 1994, 54-55; Shachar and Nalebuff 1999). Survey measures, such as whether the respondent has thought of the possibility that his or her vote might decide the election or whether or not the respondent thinks the probability of such an event is higher than “absolute zero” or “almost zero” (Blais et al. 2000) provide little incentive for truth-telling or exertion of mental effort and are imprecise measures of pivotalness.¹ Thus, the tests based on these proxies can hardly be considered as tests of the pivotal voter model (Merlo 2005).² Whether and how pivotalness factors into an individual’s turnout calculation remains largely a puzzle and so does the performance of the pivotal voter model.

¹ Harbaugh (1996), for instance, reports that in post-election surveys 10-25% of nonvoters claim that they have voted. See also Bertrand and Mullainathan (2001) for problems with survey data.

² Coate et al. (2004) test the pivotal voter model by looking at turnout in local Texas elections and considering closeness as a measure of pivotalness. However, as above, this does not provide a direct test of the model. See also Battaglini, Morton and Palfrey (2005, 21) on how simple tests of the effect of closeness on turnout are not nuanced enough as tests of the pivotal voter model.

An alternative to working with field data is to conduct laboratory experiments with paid human subjects and that is the approach we take in this paper. In the laboratory we can control both the cost of voting and the payoff to the party that wins. Using neutral language and anonymous interaction we can minimize other factors that might affect voting decisions, such as the fulfillment of “civic duty” or the avoidance of peer sanctions for non-participation. Most importantly, we can elicit the beliefs that subjects have regarding whether their voting decision will be pivotal or not, using a proper scoring rule that incentivizes truthful revelation of subjects’ beliefs. As these beliefs are fundamental to the pivotal voter model, it is important that we acquire and examine data on beliefs so as to carefully assess the extent to which agents form correct beliefs and appropriately condition their behavior on those beliefs.

The advantages of using the laboratory to test the pivotal voter model have not gone unnoticed. Several prior experimental studies have tested various aspects of the pivotal voter model, including the implications of different voting rules (plurality vs. proportional) (Schram and Sonnemans 1996a), communication, group identity and individual characteristics such as the student’s university major (Schram and Sonnemans 1996b), various comparative static predictions such as how variations in electorate sizes impact on voting decisions (Levine and Palfrey 2007) and the effects of asymmetric information (Battaglini, Morton and Palfrey 2005). However, none of these prior studies have examined the question of subjects’ subjective beliefs as to the pivotalness of their decision to vote or not, a question that lies at the heart of the pivotal voter model.

In this paper, we adopt the neutral language participation game design (Großer and Schram 2006; Schram and Sonneman 1996a, 1996b), and add to it a belief elicitation stage that

precedes the voting stage.³ In the belief elicitation stage, we ask subjects to state a subjective probability as to whether their own decision to vote or not will be decisive for the election outcome. We incentivize truthful revelation of individual beliefs using a proper scoring rule and subject earnings are determined in small part by the *ex post* accordance of their beliefs with election outcomes.⁴ In addition, we are able to study whether subjects learn over time to form correct beliefs with regard to their pivotalness in the finitely repeated election game. In sum, the study provides the first direct test of the pivotal voter model.

We find that the average participation rates are consistent with the theoretical prediction suggesting that the theoretical model works well on the aggregate level. However, our main interest is on the individual level. Here we provide some evidence that subjects are more likely to vote the higher their subjective beliefs of being pivotal – as prescribed by the pivotal voter model. Those who are certain of being pivotal are 54% more likely to participate than those who believe that their chance of being pivotal is zero. We further find that individuals learn over time to adjust their subjective probabilities of being pivotal in response to histories of voting outcomes in the direction of the true *ex post* frequency. However, although learning took place

³ See Palfrey and Rosenthal (1983) and Schram and Sonnemans (1996a) for a justification of using a participation game to represent turnout decision.

⁴ Several other experimental studies have sought to elicit subjects' subjective beliefs in environments other than the voting game that we examine. These include, but are not limited to McKelvey and Page (1990), Offerman et al. (1996), Nyarko and Schotter (2002), and Rutström and Wilcox (2004). The evidence from these studies regarding the impact of belief elicitation procedures on subject behavior is mixed. For this reason, we report data from our own control treatment without belief elicitation for the purposes of comparison.

over time, subjects consistently overestimated their probability of being pivotal and at first these predictions were quite inaccurate. Furthermore, the fit between their beliefs about decisiveness and turnout was considerably worse than the theory predicted: many subjects whose perceived pivotalness probability was higher than the cost of voting did not vote while many of those who stated a probability considerably lower than the cost of voting still decided to participate.

Overall, thus, the evidence with regards to the pivotal voter model is mixed. This should not come as a surprise. The study should not be interpreted as an attempt to “prove” or “disprove” the pivotal voter model or the rational choice theory in general. Rather, it was undertaken as an attempt to uncover whether and to what extent this theory is helpful for understanding voter turnout. The findings reported here underline the partial utility of the rationality-based arguments for such a purpose. Given this, the significant contribution of the study to the literature results from uncovering those aspects of the theory that are useful for understanding turnout decisions as well as those that are not.

Our findings are for small groups of 20 subjects. An obvious issue is whether our experimental findings “scale up” to larger electorate sizes, where the probability of being pivotal is likely to be closer to zero. We see no reason why our findings should not scale up, but acknowledge that this claim is difficult to test. Conducting controlled laboratory experiments with much larger populations is not presently feasible; internet experiments do not provide the same level of control, as one cannot rule out, e.g., communication or collusion among subjects, and survey evidence is not directly comparable to laboratory findings. On the other hand, the laboratory provides the pivotal voter theory with an *idealized* test environment – one where factors other than pivotalness (such as civic duty, or the sanction of others), have been carefully removed, and where subjects are given much more experience and information concerning

election outcomes and pivotalness than they might ordinarily encounter as voters in real elections. If the theory does not predict well in this idealized environment (with admittedly few participants) we might expect it to perform rather poorly in the less-controlled world of real elections with large numbers of participants.

Pivotal Voter Model

We consider the complete information participation game approach to modeling voting pursued by Palfrey and Rosenthal (1983). Specifically there are two teams of players of size M and N and all team members have a choice between two actions, vote (participation) or do not vote (abstention/nonparticipation). The cost of voting $c \in (0,1)$ is assumed to be the same for all agents; abstention is costless. Each member of the winning team receives a payoff benefit $B > 0$ while each member of the losing team earns a payoff of zero. The utility function is assumed to be linear, as is standard in the literature. Specifically, letting p denote the probability of casting a pivotal vote, the *net return* to voting, $R = pB - c$. Note that we abstract away from any fixed benefits to voting, e.g., the utility one gets from a “civic duty” to vote or from the avoidance of sanctions from not voting; our neutral language experimental design makes such concerns unimportant. Normalizing $B = 1$, it follows that players will rationally choose to vote whenever $p > c$, and will rationally chose to abstain if $p < c$.

The rule used to determine the outcome of voting is simple plurality. As for ties, we flipped a coin *in advance* of each election to determine which team would win in the event of a tie; the pre-announcement of the winner in the event of a tie aids in assessments of pivotalness (as described later). Given the pre-announcement of the tie-breaking rule, the setting corresponds to the “status quo” rule where there is a default winner in the event of a tie.

For our setting with $M = N > 0$ and the status quo rule, it follows from Palfrey and Rosenthal (1983) that there are no pure strategy equilibria. There may exist quasi-symmetric, totally mixed strategy equilibria where each member of the group that does not win a tie chooses to vote with probability q , defined implicitly by

$$\binom{M + N - 1}{N} q^N (1 - q)^{M-1} = c, \quad (1)$$

and members of the group that wins a tie vote with probability $1 - q$. As Palfrey and Rosenthal (1983) show, there exist values of c for which equation (1) yields either 0, 1 or 2 solutions for q .

We chose parameters for the experiment, $M = N = 10$ and $c = .18$ that are very close to the case where there is a unique, quasi-symmetric totally mixed strategy equilibrium. Our aim was to try to reduce the set of equilibria that subjects might coordinate on so as to have a more reasonable chance of predicting turnout.⁵ In the unique mixed strategy equilibrium with $M = N = 10$, $q = N / (N + M - 1) = .53$ and $1 - q = .47$ ⁶ It follows that turnout in this equilibrium involves $(2M - 1)N / (N + M - 1) = 10$ participants out of an electorate of size 20, or a turnout rate of 50 percent. While turnout is of interest to us, the primary focus of this paper is on the

⁵ There may also exist *asymmetric* equilibria, where some agents play pure strategies while others play mixed strategies, but for simplicity, we follow Levine and Palfrey (2005) and Battaglini et al. (2006) and focus on symmetric equilibria only.

⁶ The value of c needed to implement the unique mixed strategy equilibrium is .17697. Given that the smallest increment of monetary payment is .01, we chose to set $c = .18$. Technically speaking, for $c = .18$, there are two totally mixed strategy equilibria, $q_1 = .514883$ and $q_2 = .53773$, but we prefer to consider $q = .53$ as the relevant benchmark.

consistency of subjects' beliefs with their action choices. We now turn to a description of our experimental design and main hypotheses.

Experimental Design and Hypotheses

The computerized experiment was run at the _____ Experimental Laboratory of the University of _____. Subjects were recruited from the university's student population using newspaper advertisements and email. Each subject participated in only one session, and had no prior experience with our experimental setup or knowledge of our research agenda. Our experimental design involved two treatments. In the "beliefs" treatment we elicit subjects' beliefs as to whether their voting decision will be pivotal to the election outcome prior to the voting decision, while in the "control" treatment, we do not elicit beliefs. The purpose of the control treatment was to determine whether asking players about their subjective beliefs with respect to their pivotalness affected their voting behavior, e.g., made them more likely to carefully weigh the expected benefit from voting against the cost. There is conflicting evidence on the obtrusiveness of belief elicitation procedures (Offerman et al. 1996; Rutström and Wilcox 2004), and the control treatment enables us to assess whether the procedure we used impacted on voting behavior. There were three sessions of the control treatment and three sessions of the belief treatment. We first describe the control treatment; the beliefs treatment differs only in our elicitation of beliefs prior to voting decisions.

Control Treatment

In the control treatment, subjects were randomly assigned to one of two groups labeled X and Y at the start of the experimental session. We were careful to use neutral language in both treatments and avoid any context with regard to voting or elections as we did not want to cue subjects' beliefs with regard to social norms or sanction surrounding voting decisions. Subjects

were told that in each “round” of the experiment (20 rounds total), they were to decide whether to purchase a “token” or not (equivalent to casting a vote or abstaining). Purchasing a token cost them \$0.18, i.e., we set the cost of voting to $c = .18$. The payoff to each member of the winning group is \$1 while the payoff to each member of the losing group is \$0.

The experimental instructions, available at [\(author's website\)](http://(author's website)) made the payoffs to the winning team and the cost of buying a token public knowledge to all subjects. In addition, the instructions explained the plurality rule used to determine the winning group and the pre-announced tie-breaking outcome, which was to pick one team randomly each round to be the winning team in the event of a tie. Prior to the start of the experiment, subjects had to answer several quiz questions designed to test their comprehension of the rules and payoffs for the experiment. Subjects played 20 rounds of this game remaining in the same team over all rounds.⁷ They were paid their net earnings from all 20 rounds played.

The timing of moves within a round was as follows. First, the random determination as to which team will win a tie was made and announced. Second, subjects were asked to decide whether or not to purchase a token. Finally, the results of the round were revealed to all subjects. Specifically, at the end of each round, subjects were informed of the number of members of their group of 10 who purchased a token, the number of members of the other group of 10 who

⁷ We considered random re-matching of subjects into the two teams each period so as to avoid “repeated game” effects, but we decided that such a design might adversely affect subject learning, especially with regard to the probability that any individual subject is pivotal. A second consideration is that the natural field settings in which our results would be most applicable are ones that likely involve repeated interactions among the *same* individuals, e.g., members of a political party.

purchased a token and which team had won for that round. In the event of a tie, the pre-announced tie-breaking rule determined the winning group. All members of the winning group earned \$1 less the cost of purchasing a token, if they purchased a token. Similarly, all members of the losing group earned \$0 less the cost of purchasing a token, if they purchased a token.

Notice that in each round of the control treatment, subjects' *net* earnings consist of one of four possible payoffs: \$1, \$0.82, \$0 or -\$0.18; the latter negative payoff occurs when a subject buys a token and her team loses. To rule out the possibility that subjects finish the experiment with a net loss, we provided subjects with a \$6 show-up fee. As we only played 20 rounds of the voting game, the maximum loss possible was $20 \times (-.18) = \$3.60$ and subjects were informed that such losses would come out of their show-up fee. In practice, all subject payments (including the show-up payment) were greater than \$6 for both treatments. The average payoff for the 20 rounds played by subjects in the three control sessions was \$20.55 for a 90 minute experiment.

Belief Elicitation Treatment

The belief elicitation treatment differed from the control treatment in only one respect. Prior to deciding whether or not to buy a token, subjects were asked to report their subjective belief as to whether their decision to buy a token would be decisive (pivotal) or not.⁸ To aid subjects in formulating this belief, the conditions under which their decision to buy or not buy a token would be decisive were carefully explained in the experimental instructions. The decisiveness conditions made use of the fact that one group was randomly selected at the start of each round to be the winning group in the event of a tie.

⁸ For current purposes, we consider the terms “decisive” and “pivotal” as synonyms. In the instructions we used the term “decisive” in order to make the concept easier to understand for the subjects. Subjects were given a precise working definition of “decisive” (see below).

The timing of moves within a round was as in the control treatment, but with the addition of the belief elicitation stage. Specifically, the timing was as follows. First, the random determination as to which team would win a tie was announced. Second, subjects then stated their subjective belief as to whether their decision to purchase a token would be decisive. Third, subjects were asked to decide whether or not to purchase a token. Finally, the results of the round were revealed as in the control treatment. The information revealed at the end of each round included the same information that was revealed at the end of a control session, and additionally, subjects were reminded of their stated belief and whether their token purchase decision was decisive or not for the outcome of the round. The latter information was intended to provide subjects with the feedback necessary to better align their decisiveness beliefs with actual outcomes. It is perhaps useful to quote the instructions with regard to the conditions under which an individual subject's token purchase decision is decisive:

“You are decisive under any of the following conditions.

Suppose that group X wins a tie.

1. If there is a tie then everyone in group X who *bought* a token is decisive.
2. If there is a tie then everyone in group Y who *did not buy* a token is decisive.
3. If group X loses by one token, then everyone in group X who *did not buy* a token is decisive.
4. If group Y wins by one token, then everyone in group Y who *bought* a token is decisive.

Suppose instead that Y wins a tie.

1. If there is a tie then everyone in group Y who *bought* a token is decisive.
2. If there is a tie then everyone in group X who *did not buy* a token is decisive.
3. If group Y loses by one token, then everyone in group Y who *did not buy* a token is decisive.
4. If group X wins by one token, then everyone in group X who *bought* a token is decisive.”

To make it incentive compatible for subjects to report their true beliefs regarding decisiveness, we make use of a proper scoring rule and give subjects a small payment according

to the accuracy of their stated beliefs. Specifically, we make use of the quadratic scoring rule originally developed by Brier (1950) for weather forecasting but more recently adopted by many experimentalists (McKelvey and Page 1990; Nyarko and Schotter 2002; Offerman et al. 1996 among others). Suppose a subject reports the subjective probability p that she will be decisive. *Ex post*, when the election results are determined, either she is decisive or not. Let I_d be an indicator function that takes on the value “1” if the subject is decisive and “0” otherwise. The payoff we give to subjects for their stated belief each round is $\pi(p) = .10[1 - p - I_d]^2$. That is, the maximum subjects can earn for a correct guess is \$0.10 and this amount diminishes quadratically as the guess deviates from the actual outcome, down to \$0.00. Theoretically, the quadratic scoring rule induces a risk-neutral agent to report her true, subjective belief with regard to the binary event, in our case, being decisive in the participation game (Winkler and Murphy 1968; see also the discussion in Camerer 1995, 592-3). In setting the payoff for the decisiveness prediction, we followed Nyarko and Schotter (2002) in making this payoff small with respect to the payoff of winning an election (which was \$1). A concern is that when subjects earn payments for both their action choices and for the accuracy of their stated beliefs, they might state beliefs in a strategic fashion that provided them with insurance against election outcomes. By keeping the payment for belief accuracy small, we sought to minimize such strategic behavior.

Aside from elicitation of beliefs before voting decisions, there were no differences between the two treatments. Subjects in the belief-elicitation treatments answered several additional quiz questions that tested their comprehension of the decisiveness rules and payoff possibilities in the belief treatment. They earned slightly more on average, \$21.75 than subjects in the control treatment, but the differences are easily accounted for by the additional payments subjects received in the former treatment for the accuracy of their beliefs.

Our main interest in the belief elicitation treatment is to assess the extent to which subjects choose to vote when their decisiveness beliefs exceed the cost of voting, $p > c = .18$, and abstain from voting otherwise. We are also interested in whether subjects learn over time to adjust their beliefs toward the actual frequency of being decisive and whether our belief elicitation procedure has an impact on decision-making, perhaps by making subjects think harder about the rationality of choosing to buy a token.

Results

We report results from six sessions – three belief elicitation treatment sessions and three control treatment sessions. Each session involved 20 subjects who made decisions in 20 rounds. Thus, there are 1200 participation (or voting) decisions from the belief elicitation treatment and the same number from the control treatment, i.e. a total of 2400 decisions.

Aggregate results

Figure 1 and Table 1 summarize turnout using 5-round averages for each of the six sessions, labeled beliefs or no-beliefs sessions 1, 2, 3. The average turnout in rounds 1-5 is around 50%. It, thus, corresponds well with the theoretical prediction (see above). However, in all sessions, the turnout levels drop off over time. The turnout levels and trends across treatments are similar enough to suggest that our belief elicitation procedure was not obtrusive. To be sure, turnout is slightly higher when beliefs were elicited: the average turnout is 45% in the belief treatment and 39% in the control treatment. Further, the turnout level drops off slightly more sharply in the control sample. However, the differences across the two treatments are not large; using a non-parametric Mann-Whitney test (results in the final column of Table 1), the null hypothesis of no difference in turnout rates between treatments can be rejected only in rounds 15-20 at the .05 level of significance. In one of the no-belief sessions, session 2, one group became dominant,

i.e. the same group was winning in all rounds. This may account for lower participation rate in this session (see Figure 1). In all other sessions no dominant group emerged. Additionally, the share of decisive games was very similar across treatments: in the belief sample, 14 out of the 60 games resulted in a decisive participation, while in the control treatment the ratio was 12 out of 60. We will return to the issue of the potential obtrusiveness of the experimental treatment below. To foreshadow the conclusion, we find no significant differences in the results across treatments indicating the unobtrusiveness of the belief elicitation procedure.

Individual level results

The crucial independent variable in this study is the subjective decisiveness probability. Subjects could state a probability with an accuracy of up to three decimal places. In all sessions, “0” and “.5” were modal values, though many other values were chosen. The mean subjective probability that an individual is decisive is rather high: .33. It varies slightly by session, equal to .29 for the first and the third belief sessions and .41 for the second. Figure 2 shows frequency and cumulative frequency distributions for the subjective decisiveness probabilities by session averaged over the first and last 10 rounds. As the graphs illustrate, subjects’ decisiveness probabilities in the first 10 rounds are spread more uniformly over the interval $[0,1]$, as compared with the last 10 rounds where the distribution is more skewed to the left of the interval. This is also reflected in the cumulative frequency distributions, which for rounds 11-20 nearly always lie above those for rounds 1-10. Figure 3 presents the mean subjective decisiveness probabilities over time. The first session shows clear convergence in beliefs across subjects over time: the slope of the fitted regression line is negative and significant, and the standard deviation bounds around the mean narrow considerably towards the end of the series. A similar convergence is less evident in the other two sessions, although in both cases the fitted slopes are negative.

On average, 63% of subjects across all 3 belief elicitation sessions stated a probability of being decisive that was higher than .18. Recall that $c = .18$; thus, the decisiveness probability of .18 serves as the theoretical cutpoint for participation. These subjective probabilities of being decisive can be compared to the actual probabilities, or the frequencies of past decisiveness. The actual mean frequency of decisiveness (all 20 rounds) averages out to be .131 across all 3 belief sessions (see also Table 2 below); the mean decisiveness frequency is highest in the second (.21), lowest in the first (.055) and intermediate in the third session (.128).⁹ The difference between the historical average objective frequency of decisiveness and the subjective frequency of decisiveness are rather substantial. Figure 4 illustrates, by sessions, how this difference between objective average historical decisiveness and the average subjective decisiveness decreases over time.¹⁰ The convergence is especially visible in the case of the first session (solid line) where in the last two rounds the objective and average subjective probabilities are equal. This indicates that individuals can learn over time to adjust their subjective probabilities of being pivotal in response to histories of voting outcomes in the direction of the true *ex post* frequency of decisiveness. The positive values of the series in Figure 4 indicate that subjects are almost without an exception *overestimating* the probability.

Another, complementary means of assessing the accuracy of subjects' forecasts is to examine a simple linear transformation of the quadratic scoring rule that was used to determine

⁹ In theory, in the mixed strategy equilibrium, the frequency of decisiveness would average .18.

¹⁰ The average historical decisiveness at the start of round t is the average frequency with which subjects have been decisive in all prior rounds $t=1, \dots, t-1$. Figure 4 plots the average difference between subjects' stated subjective probability of decisiveness for round t and the average historical decisiveness at the start of round t , beginning with round 2.

subjects' forecast earnings known as the *probability score* (or “*Brier score*” - see, e.g. Yates 1994). The probability score for player i in round t is given by:

$$ps_t^i = (p_t^i - d_t^i)^2, \quad \text{where } d_t^i = \begin{cases} 1 & \text{if } i \text{ was decisive at round } t, \\ 0, & \text{otherwise.} \end{cases}$$

and p_t^i is the stated decisiveness probability of subject i at the start of round t .

As one benchmark for this statistic, notice that if a subject never has a clue as to her probability of being decisive, and therefore *sensibly* chooses to state the uniform prior, $p = .5$ in every round (perhaps a big assumption!) her probability score would be .25 in every round. At the other extreme is a benchmark provided by the fictional “base rate judge”, who (inexplicably) knows the sample average frequency of decisiveness, \bar{d} (over all subjects, all rounds) and states $p = \bar{d}$ in every round. Yates (1994) shows using a simple decomposition that the mean probability score for a base rate judge, $\overline{ps}^b = \bar{d}(1 - \bar{d})$, which is just the sample variance of the decisiveness index.¹¹

In Table 2 we report mean probability scores over all subjects for each of the three sessions in which beliefs were elicited, using data from the first ten, the second ten and from all 20 rounds. We also report the sample mean decisiveness \bar{d} , and \overline{ps}^b . We see that in two of three sessions (1 and 3) subjects' probability scores average around .25 in the first 10 rounds, and drop below this level in the second 10 rounds. However in one session (2), the probability scores initially average above .25 in the first 10 rounds and this average increases over the second 10 rounds! Averaging over all 3 sessions, there is a slight decline in the mean probability score from .253 in the first 10 rounds to .207 in the second 10 rounds, suggesting that, on average across sessions, there is *some* improvement over time in the accuracy of subject's probability

¹¹ A perfect probability score of 0 will generally lie below the score of a base-rate judge.

assessments. Notice too, that while subjects' probability scores are well above those of a base rate judge, as given by \bar{p}^b , consistent with overestimation of pivotalness as discussed above, the variation in probability scores across sessions is perfectly correlated with the variation in the probability scores of the base rate judge, with the highest probability scores observed in session 2, the lowest in session 1 and session 3 having intermediate probability scores.

In addition to assessing the accuracy of beliefs, we can also examine the *payoff efficiency* of subjects' decisions in the belief treatment relative to the theoretical, rational choice benchmark. Specifically, we calculated the payoffs subjects *would have earned* had they played strict best responses to their subjective probabilities of pivotalness in each round, i.e., if they had behaved according to rational choice theory, choosing to vote whenever their stated p was greater than .18 and choosing to abstain whenever their stated p was less than .18 (no instances of $p = .18$ were found in the data). This hypothetical exercise, of course, leads to different election outcomes and payoffs than are found in the actual data. In the event of ties, we used the actual, pre-announced tie-breaking rule for the round (which was announced *before* subjects submitted their subjective probabilities).

Table 3 reports the ratio of actual payoffs to hypothetical, "rational choice" benchmark payoffs for each of the three belief sessions over the first 10, the last 10 and all 20 rounds.¹² It also reports the average difference in actual and hypothetical payoffs. We see that, with a single exception, (the last 10 rounds of belief session 1) that subjects were earning slightly *higher* payoffs than they would have had they played according to the rational choice benchmark prediction. The average differences between actual and hypothetical payoffs are, in all instances,

¹² For simplicity, both the actual and hypothetical payoffs used in these ratios did not include the small payoff component that subjects earned for the accuracy of their stated beliefs.

quite small, just a few cents. Of course, had subjects played according to the rational choice prediction, their subjective beliefs might have evolved differently and so our comparison is not exact in that dimension. Still, this finding suggests that, while subjects were not playing crisp best responses to their stated beliefs (more on this below), they appear to have been no worse off as a result, and indeed may have been slightly better off!

Multivariate analyses

In order to further understand the effect of subjective beliefs of pivotalness on the likelihood of buying a token we have conducted a number of multivariate probit regressions. As individual decisions within sessions are not entirely independent, we have clustered the standard errors on subjects in all analyses and included session dummies. The results are presented in Table 4.

Model 1 estimates the effect of the stated beliefs of being pivotal (continuous variable) on the decision to vote (binary variable) while Model 2 replicates the same analysis using a dummy variable coded “1” for those who stated a probability of being pivotal higher than .18 in order to test the exact predictions of the theory.

Both models include several controls. First, we control for whether the group of which the subject is a member will win in the event of a tie. This variable might also be thought of as proxying for a pre-election poll announcing a lead to one candidate. The pivotal voter model predicts lower turnout for the “advantaged” group (see fn. 6; Levine and Palfrey 2007). Further, since we ran several rounds of “elections” and the group members stayed the same across rounds, we also control for various history effects. These include (1) whether a given subject was pivotal in the last round, (2) whether the subject bought a token in the last round, (3) whether the subject’s group won last round, (4) the number of tokens bought by the subject’s group in the last

round, (5) the subject's earnings from the last round, and (6) whether there was a tie in the last round. We also control for session effects using session dummies and for the round number.

The results of Model 1 show a strong effect of the stated probability of being decisive on the probability of buying a token. Substantively, the predicted probability of buying a token is .37 when the stated probability of being pivotal is 0 (i.e. at its minimum) and .58 when it is 1 (at its maximum), holding other variables at their mean (for continuous variables) or median (for categorical variables). In Model 2, however, the effect of the decisiveness probability is not statistically significant. Substantively the effect is also marginal: the predicted probability for buying a token is .4 when the stated probability of being pivotal is higher than .18 and only slightly higher – .47 – when it is lower than .18, all other variables at their mean or median.

There is, thus, some evidence that the subjective probability of being pivotal plays a significant role in people's decision to participate: the higher the subjective probability the greater the likelihood of buying a token. The result is not, however, as crisp as the theory would predict: a subjective probability of .18 does not function as a clear cutpoint for the decision to participate. Figure 5 further illustrates this point. The left panel of this figure shows the percentage of subjects who bought a token (chose to participate) after having stated a probability of being pivotal that was less than .18 while the right panel shows the percentage of subjects who bought a token after having stated a probability of being pivotal that was greater than or equal to .18. The figures show average percentages over rounds 1-2 and 19-20 of each session. If players were playing according to the crisp cutpoint prediction of the theory, the percentages in the left panel would all be zero while those in the right panel would all be 100. Notice however, that in 2 of the 3 belief sessions, the participation rate when $p < .18$ decreases from rounds 1-2 to the final rounds 19-20, while in all 3 belief sessions, the participation rate increases when $p \geq .18$ from

rounds 1-2 to the final rounds 19-20; the latter finding suggests that subjects may be learning to behave in a manner consistent with the theoretical cutpoint prediction.¹³ Further, Figure 6 plots average beliefs of being decisive over time for participants and non-participants. As is evident from these graphs, although the decisiveness probabilities of participants are usually higher than those of non-participants, there does not appear to be a clear average cutpoint for participation. In sum, there is only weak support for the specific prediction of the theory. Few participants use the exact deterministic cutpoint strategy predicted by the theory. However, there is evidence that subjects' behavior tends towards the theoretical prediction with higher subjective probabilities increasing the likelihood of participation. Furthermore, as the results presented above, subjects' payoff efficiency is already approximately equal to that of a rational choice voter.

Additional findings

In addition to the main findings, some of the variables measuring the effects of history or past behavior are also significantly related to the decision to participate. First, the round has a significant negative effect on the probability of buying a token: all else equal, subjects were less likely to buy a token in later than in earlier rounds. This may indicate a certain learning effect, for example in terms of cumulative disappointment in low payoffs from buying a token, or the emergence of a free rider problem (see Bendor, Diermeier and Ting 2003; Kanazawa 2000 for learning effects). This result also reflects the observation that turnout declines when democracies

¹³ Alternatively, this kind of noisy cut-point behavior might be well-approximated by a quantal response equilibrium (McKelvey and Palfrey 1995), an important generalization of Nash equilibrium that allows for noisy best responses. However, our focus here is not so much on the Nash equilibrium predictions (i.e., with regard to turnout) as it is on the individual rationality of voting decisions given a subjects' own stated, subjective belief of being pivotal.

mature, i.e. as a result of repeated elections (Kostadinova 2003). Second, subjects are more likely to participate when they have participated before. This result reflects the argument about the “habitual voter” (Gerber, Green and Shachar 2003; Plutzer 2002) made in a previous empirical literature on turnout. Further, the participation rate of one’s group members also significantly influences individual’s decision to buy a token: a subject is less likely to participate if the general participation rate in his or her group was high indicating the emergence of a free rider problem. Past group success or failure does not play a significant role in decisions to participate.

Models 1 and 2 also allow us to test a further implication of the theory. Recall that in the unique mixed strategy equilibrium the turnout probabilities for the two groups are not the same (see fn. 6). Rather, we predicted an underdog effect (Levine and Palfrey 2007), where the group winning a tie should have higher turnout than the disadvantaged group. The “Group wins tie” variable in Models 1 and 2 captures this relationship. Contrary to the theory, the finding indicates that the tie breaking rule acts as a coordination device for voters, mobilizing (rather than demobilizing) them behind a “leading” candidate. Substantively, in both models, the predicted probability of buying a token is .34 for a member of the group that is announced to win a tie compared to .25 for a member of the other group. Figure 7 further illustrates this effect. It presents the average turnout in all six sessions for the advantaged and disadvantaged groups. In all sessions, the turnout is higher for the former, the difference being statistically significant (using a non-parametric Mann-Whitney test) in all but the second belief session. Although our theory does not predict such a result, it is consistent with the *bandwagon effect* reported in several voting studies, according to which voters favor a party that is doing well in the polls (McAllister and Studlar 1991). In the current context, this finding may have occurred as an artifact of our particular experimental design, i.e., the use of *ex ante* coin toss (or status quo) rule

for breaking ties and fixed group membership. An alternative, yet strategically equivalent design might declare in advance that one group always wins a tie and employ random re-matching of subjects into groups as opposed to our fixed group membership design. If the bandwagon effect disappears, then we might attribute the current result to our design.¹⁴

In two additional models – Model 3 and 4 in Table 4, we added the variable measuring the average historical frequency of being decisive. Adding this variable does little to diminish the effect of the subjective probabilities of being decisive. Rather, the objective frequencies have no statistically significant effect on turnout and the coefficient has a wrong sign, while the effect of subjective beliefs remains still significant and in the predicted direction. This underlines the importance of subjective beliefs of being pivotal in turnout decisions and challenges the use of some objective measures of this probability when testing the pivotal voter model such as closeness of an election. As Figure 4 demonstrated, although over time the subjective probability of being pivotal tends towards the actual frequency, the differences can be substantial.

Testing the obtrusiveness of the belief elicitation procedure

Models 5 and 6 in Table 4 replicate Model 3 with data from the control treatment and from all sessions combined respectively. The goal here is to determine whether subjects behaved significantly differently when beliefs were elicited compared to the control group. Most importantly, the dummy variable differentiating between treatment and control sessions (variable name *Belief elicitation*) in Model 6 is not statistically significant. This allows us to conclude that there are no significant differences in the behavior of subjects across treatments and that our belief elicitation procedure was not obtrusive in terms of making subjects more aware of the

¹⁴ A random re-matching design, however, might make it more difficult for subjects to assess or adjust their subjective probability of decisiveness. We leave this exercise to future research.

rationality of participating. Furthermore, Model 5 produces roughly similar results as Model 3 in Table 4. As above, we find that the tie breaking rule influences turnout, prior participation increases while high level of group participation depresses the likelihood of current participation, and turnout decreases significantly over time, while the effect of the historical frequency of decisiveness falls short of statistical significance. These similarities across treatments further suggest that the belief elicitation was unobtrusive in terms of influencing subjects' decision making. There are also some differences however: both prior earnings and history of ties significantly and positively influence turnout. These history effects are likely influencing one's subjective probability of being pivotal, and as such, the relationships may reflect some of the effects otherwise captured by the subjective beliefs.

Conclusions

The pivotal voter model that builds on Downs' (1957) rational choice theory of turnout is the most intuitive, yet also the most controversial formal theory in political science. Empirical tests of the theory to date have mostly relied on proxies or have been partial. The concept around which much of the controversy revolves – the probability of being pivotal – has received the least empirical attention. Indeed, pivotalness is simply inferred from the closeness of the election, and the individual level calculus of voting is never unpacked. This study has provided the first direct test of the pivotal voter model, with a specific focus on whether and how voters' beliefs of being pivotal factor into their voting decisions. We used a laboratory setting that allows manipulating model parameters and elicited voters' beliefs about pivotalness.

Importantly, we determined that the belief elicitation procedure did not prompt subjects to alter their turnout behavior. On the aggregate level, we find support for the pivotal voter model with aggregate turnout rates corresponding well with the theoretical predictions.

The main focus of this study, however, is on the individual level behavior, where we find mixed support for the theory. Most importantly, we find evidence that those who believe that their participation will be pivotal for the outcome of a game (an election), are more likely to participate. This relationship is strong and robust, but it is not deterministic – there is no single cutpoint strategy of participation that applies to all voters as predicted by the theory. Those who state a probability of being pivotal higher than the cost of voting participate with a higher likelihood than those who state a lower probability, but many voters stating much higher probabilities of being pivotal still chose to abstain and those who stated a lower probability chose to participate. That is, there is a stochastic element in voter choice. At the same time, the subjects' payoff efficiency is already approximately equal to that of a rational choice voter, so the payoff incentives might be insufficient to move subjects closer to the threshold prediction of the theory.

Additionally, we have found that subjective beliefs are more important for the participation (turnout) decision than the actual frequencies of being decisive. Indeed, the subjective probabilities tend to be considerably higher than the actual ones – undermining the common view that closeness is a useful proxy for pivotalness in testing rational models of turnout. Yet, we also find that, on average, beliefs become somewhat more accurate over time, indicating a learning effect in the turnout decision. The overestimation of pivotalness may, thus, provide a solution to the paradox of voter turnout – voting happens because people systematically think that their vote counts more than it actually does, though this overestimation declines with experience. The across-time learning process would then also account for the declining turnout in almost all democracies across the world.

References

- Aldrich, John, H. 1993. "Rational Choice and Turnout." *American Journal of Political Science* 37(1): 246-278.
- Battaglini, Marco, Rebecca Morton and Thomas Palfrey. 2006. "The Swing Voter's Curse in the Laboratory." Working Paper, Princeton University, California Institute of Technology and New York University.
- Bendor, Jonathan, Daniel Diermeier, and Michael Ting. 2003. "A Behavioral Model of Turnout." *American Political Science Review* 97(2): 261-280.
- Bertrand, Marianne and Sendhil Mullainathan. 2001. "Do People Mean What They Say? Implications for Subjective Survey Data" *American Economic Review* 91(2): 67-72.
- Blais, André, Robert Young and Miriam Lapp. 2000. "The Calculus of Voting: An Empirical Test." *European Journal of Political Research* 37: 181-201.
- Brier, Glenn W. 1950. "Verification of Forecasts Expressed in Terms of Probability," *Monthly Weather Review* 78:1-3.
- Camerer, Colin. 1995. "Individual Decision Making" in John H. Kagel and Alvin E. Roth, Eds. *The Handbook of Experimental Economics*, Princeton: Princeton University Press.
- Coate, Stephen, Michael Conlin and Andrea Moro. 2004. "The Performance of the Pivotal-Voter Model in Small-Scale Elections: Evidence from Texas Liquor Referenda." Unpublished manuscript. Cornell University.
- Cyr, A. Bruce. 1975. "The Calculus of Voting Reconsidered." *Public Opinion Quarterly* 39(1): 19-38.
- Dhillon, Amrita and Susana Peralta. 2002 "Economic Theories of Voter Turnout." *Economic Journal* 112: F332-F352.

Downs, Anthony. 1957. *An Economic Theory of Democracy*. New York: Harper & Row.

Ferejohn, John and Morris P. Fiorina. 1975. "Closeness Counts Only in Horseshoes and Dancing." *American Political Science Review* 69: 920-925.

Fiorina, Morris. 1990. "Information and Rationality in Elections." In *Information and Democratic Processes*, ed. John Ferejohn and James Kuklinski. Urbana: University of Illinois Press.

Foster, Carroll B. 1984. "The Performance of Rational Voter Models in Recent Presidential Elections." *American Political Science Review* 78(3): 678-690.

Friedman, Jeffrey (ed.). 1995. *The Rational Choice Controversy*. New Haven: Yale University Press.

Gerber, Alan S., Donald P. Green and Ron Shachar. 2003. "Voting May be Habit-Forming: Evidence from a Randomized Field Experiment." *American Journal of Political Science* 47(3): 540-550.

Green, Donald P. and Ian Shapiro. 1994. *Pathologies of Rational Choice Theory*. New Haven: Yale University Press.

Großer, Jens and Arthur Schram. 2006. "Neighborhood Information Exchange and Voter Participation: An Experimental Study." *American Political Science Review* 100(2): 235-248.

Hansen, Steven, Thomas Palfrey and Howard Rosenthal. 1987. "The Relationship Between Constituency Size and Turnout: Using Game Theory to Estimate the Cost of Voting." *Public Choice* 52: 15-33.

Harbaugh, William T. 1996. "If People Vote Because They Like To, Then Why Do So Many of Them Lie?" *Public Choice* 89: 63-76

- Kanazawa, Satoshi. 2000. "A New Solution to the Collective Action Problem: The Paradox of Voter Turnout." *American Sociological Review* 65(3): 433-442.
- Kostadinova, Tatiana. 2003. "Voter Turnout Dynamics in Post-Communist Europe." *European Journal of Political Research* 42: 741-759.
- Ledyard, John O. 1984. "The Pure Theory of Large Two-Candidate Elections" *Public Choice* 52(1): 15-33.
- Levine, David K. and Thomas R. Palfrey. 2007. "The Paradox of Voter Participation? A Laboratory Study." *American Political Science Review* 101(1): 143-158.
- Matusaka, John G. and Filip Palda. 1993. "The Downsian Voter Meets the Ecological Fallacy." *Public Choice* 77: 855-878.
- McAllister, Ian and Donley T. Studlar. 1991. "Bandwagon, Underdog, or Projection? Opinion Polls and Electoral Choice in Britain, 1979-1987." *Journal of Politics* 53(3): 720-741.
- McKelvey, Richard D. and Talbot Page. 1990. "Public and Private Information: An Experimental Study of Information Pooling," *Econometrica*, 58: 1321-39.
- McKelvey, Richard D. and Thomas R. Palfrey. 1995. "Quantal Response Equilibria for Normal form Games." *Games and Economic Behavior* 10: 6-38.
- Merlo, Antonio. 2005. "Whither Political Economy? Theories, Facts and Issues." Paper presented at the World Congress of the Econometric Society, London, August 19-24.
- Nyarko, Yaw and Andrew Schotter. 2002. "An Experimental Study of Belief Learning Using Elicited Beliefs," *Econometrica*, 70: 971-1005.
- Offerman, Theo, Joep Sonnemans and Arthur Schram, 1996. "Value Orientations, Expectations and Voluntary Contributions in Public Goods." *Economic Journal*, 106: 817-45.

- Palfrey, Thomas R. and Howard Rosenthal. 1983. "A Strategic Calculus of Voting." *Public Choice* 41: 7-53.
- Palfrey, Thomas R. and Howard Rosenthal. 1985. "Voter Participation and Strategic Uncertainty." *American Political Science Review* 79(1): 62-78.
- Plutzer, Eric. 2002. "Becoming a Habitual Voter: Inertia, Resources, and Growth in Young Adulthood." *American Political Science Review* 96(1): 41-56.
- Riker, William and Peter Ordeshook. 1968. "A Theory of the Calculus of Voting." *American Political Science Review* 79(1): 62-78.
- Rutström, E. Elisabet and Nathaniel T. Wilcox 2004. "Learning and Belief Elicitation: Observer Effects." Unpublished manuscript. University of Houston.
- Schram, Arthur and Joep Sonnemans. 1996a. "Voter Turnout as a Participation Game: An Experimental Investigation." *International Journal of Game Theory* 25: 385-406.
- Schram, Arthur and Joep Sonnemans. 1996b. "Why People Vote: Experimental Evidence." *Journal of Economic Psychology* 17: 417-442.
- Shachar, Ron and Barry Nalebuff. 1999. "Follow the Leader: Theory and Evidence on Political Opinion." *American Economic Review* 89(3): 525-547.
- Tullock, Gordon. 1967. *Towards the Mathematics of Politics*. Ann Arbor: University of Michigan Press.
- Winkler, R.L. and A.H. Murphy. 1968. "'Good' probability assessors." *Journal of Applied Meteorology* 7: 751-8.
- Yates J. Frank. 1994. "Subjective Probability Accuracy Analysis," in G. Wright and P. Ayton eds., *Subjective Probability*, New York: John Wiley & Sons.

Table 1: Turnout rates, 5-round averages, 3 belief and 3 no-belief sessions

Rounds	Beliefs 1	Beliefs 2	Beliefs 3	No Beliefs 1	No Beliefs 2	No Beliefs 3	MW null hypo
							p-value
1-5	0.49	0.50	0.47	0.53	0.40	0.42	0.80
6-10	0.51	0.51	0.42	0.39	0.47	0.41	0.10
11-15	0.47	0.42	0.40	0.38	0.26	0.44	0.20
16-20	0.35	0.41	0.41	0.31	0.28	0.27	0.05

Table 2: Mean (st. dev.) in decisiveness (pivotalness) probability scores, 3 belief sessions

Session	Periods 1-10	Periods 11-20	Periods 1-20	\bar{d}	\bar{ps}^b
1	.245 (.262)	.082 (.167)	.163 (.234)	.055	.052
2	.264 (.259)	.323 (.317)	.293 (.290)	.210	.166
3	.251 (.312)	.217(.304)	.234 (.308)	.128	.111
All 3	.253 (.278)	.207 (.288)	.230 (.284)	.131	.114

Table 3: Comparison of Actual to Hypothetical “Rational Choice” Payoffs in the 3 Belief Elicitation Sessions

Ratios of Actual to Hypothetical Payoffs			
Session	Periods 1-10	Periods 11-20	Periods 1-20
1	1.112	0.952	1.024
2	1.149	1.168	1.159
3	1.059	1.070	1.064
All 3	1.106	1.056	1.080

Average Difference Between Actual and Hypothetical Payoffs			
Session	Periods 1-10	Periods 11-20	Periods 1-20
1	0.0414	-0.0216	0.0099
2	0.0531	0.0612	0.05715
3	0.0234	0.0279	0.02565
All 3	0.0393	0.0225	0.0309

Table 4: Probit analysis of the subjective probability of being decisive (pivotal) on voting

	Experimental group (beliefs elicited)				Control group	All groups
	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
Subjective Pr(Decisive)	.518** (.201)		.544** (.203)			
Subjective Pr(Decisive) >.18		.157 (.132)		.164 (.141)		
Beliefs elicited						.113 (.179)
Historical frequency of decisiveness			-.472 (.458)	-.357 (.453)	1.219 (.935)	.056 (.398)
Group wins tie	.250** (.101)	.252** (.100)	.249** (.101)	.252** (.100)	.549** (.129)	.386** (.083)
Decisive <i>t</i> -1	.109 (.132)	.150 (.130)	.188 (.161)	.212 (.161)	.218 (.143)	.247** (.111)
Participate <i>t</i> -1	.832** (.147)	.805** (.147)	.831** (.147)	.803** (.146)	.861** (.152)	.814** (.104)
Win <i>t</i> -1	-.001 (.092)	-.005 (.090)	-.006 (.091)	-.008 (.089)		.091 (.083)
Number of group tokens <i>t</i> -1	-.097** (.029)	-.093** (.029)	-.096** (.029)	-.092** (.029)	-.112** (.026)	-.099*** (.019)
Earnings <i>t</i> -1	.050 (.114)	.026 (.114)	.039 (.112)	.016 (.112)	.322** (.142)	.091 (.088)
Tie <i>t</i> -1	-.117 (.168)	-.098 (.174)	-.129 (.170)	-.106 (.174)	.581** (.185)	.079 (.139)
Constant	-.291 (.233)	-.215 (.243)	-.275 (.233)	-.202 (.243)	-.341* (.172)	-.291 (.162)
X^2	69.41**	54.69**	69.16**	54.32**	123.56**	153.04***
Pseudo R^2	.09	.08	.09	.08	.12	.1
N	1140	1140	1140	1140	1140	2280

Note: Dependent variable is whether or not a token was bought. Table entries are probit coefficients with robust standard errors in parentheses. Standard errors are clustered on subject. * $p \leq .05$, ** $p \leq .01$. Session dummies and a trend variable for rounds are included but not reported. In Model 4 Win *t*-1 is dropped due to co-linearity.

Figure 1: Turnout rates for all sessions, 5-round averages

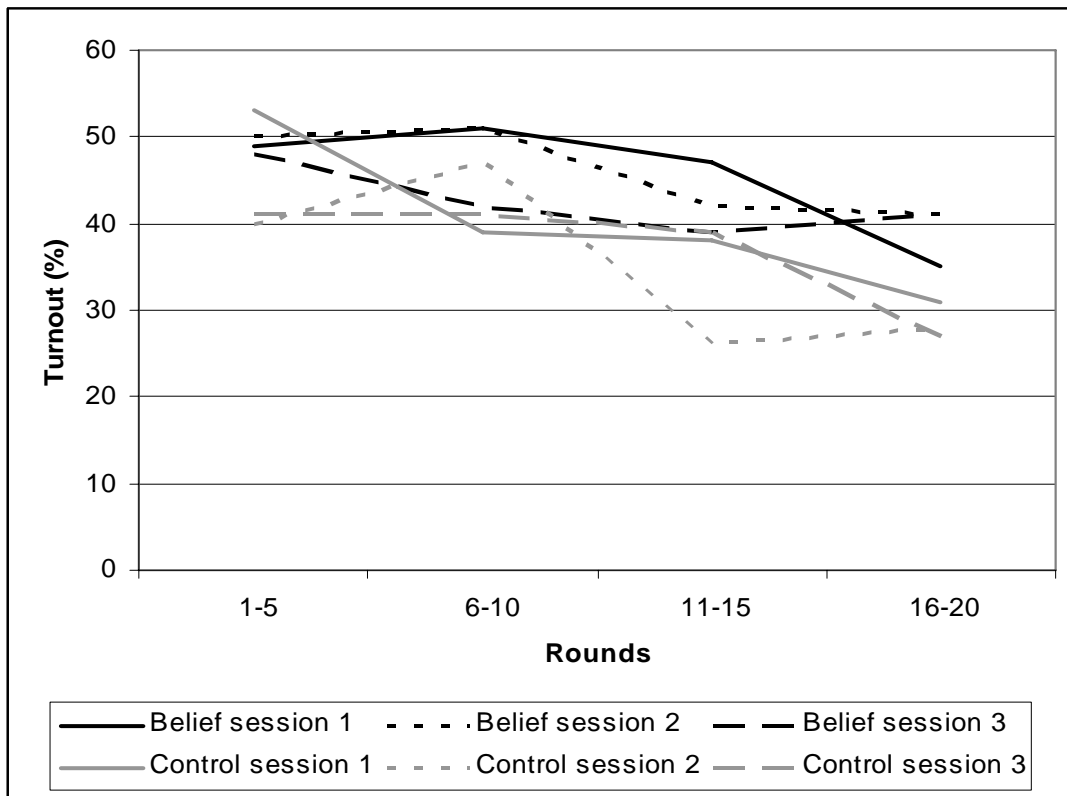


Figure 2: Frequency (left column) and cumulative frequency distribution (right column) of subjective decisiveness (pivotalness) probabilities by session, averages over round 1-10, 11-20.




Figure 3: Mean decisiveness (pivotalness) probabilities over time by session


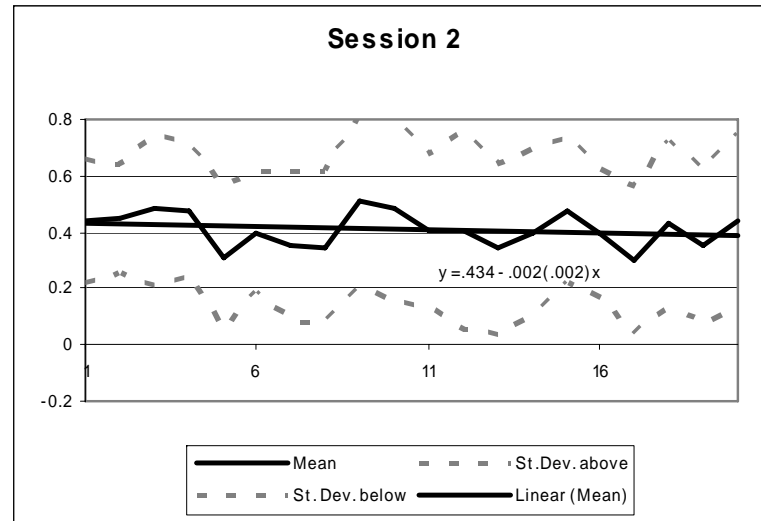
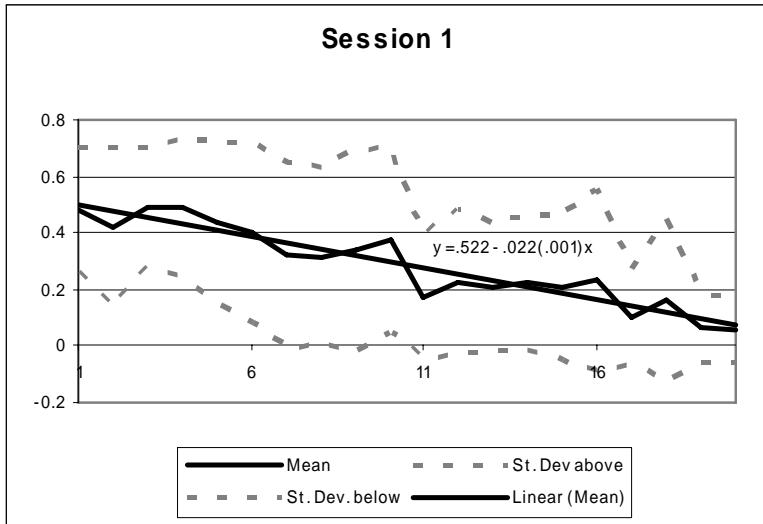


Figure 4: The difference in the subjective probability and average historical frequency of being decisive (pivotal)

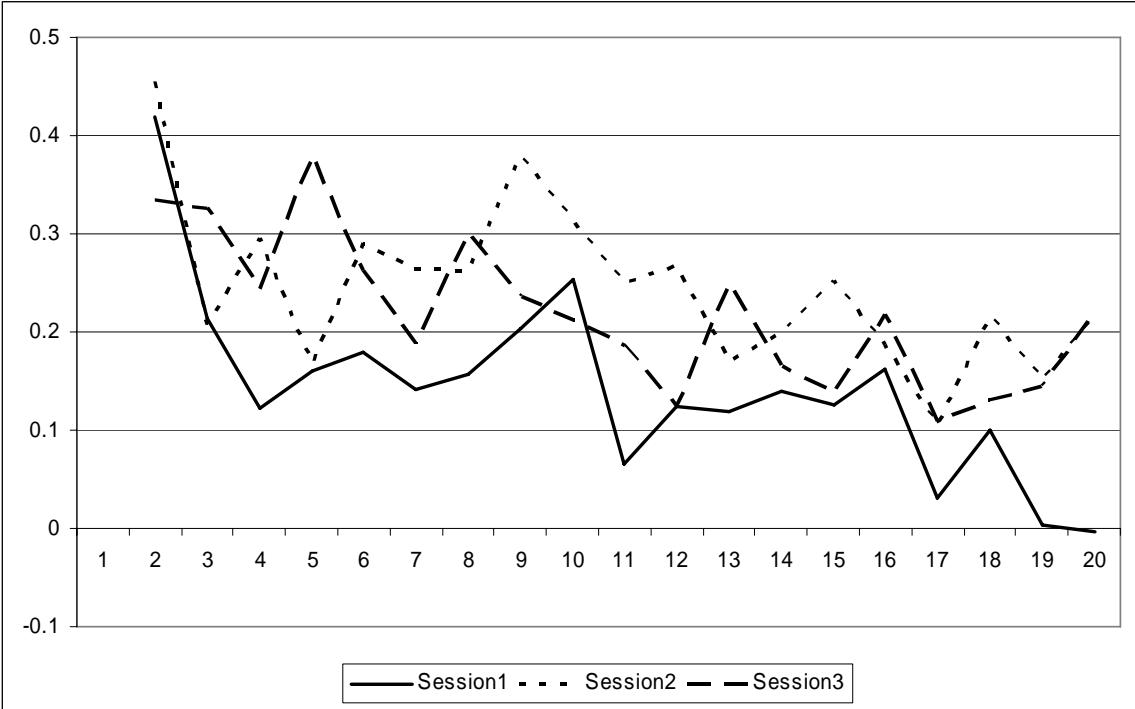


Figure 5: Participation rate by decisiveness (pivotalness) probability

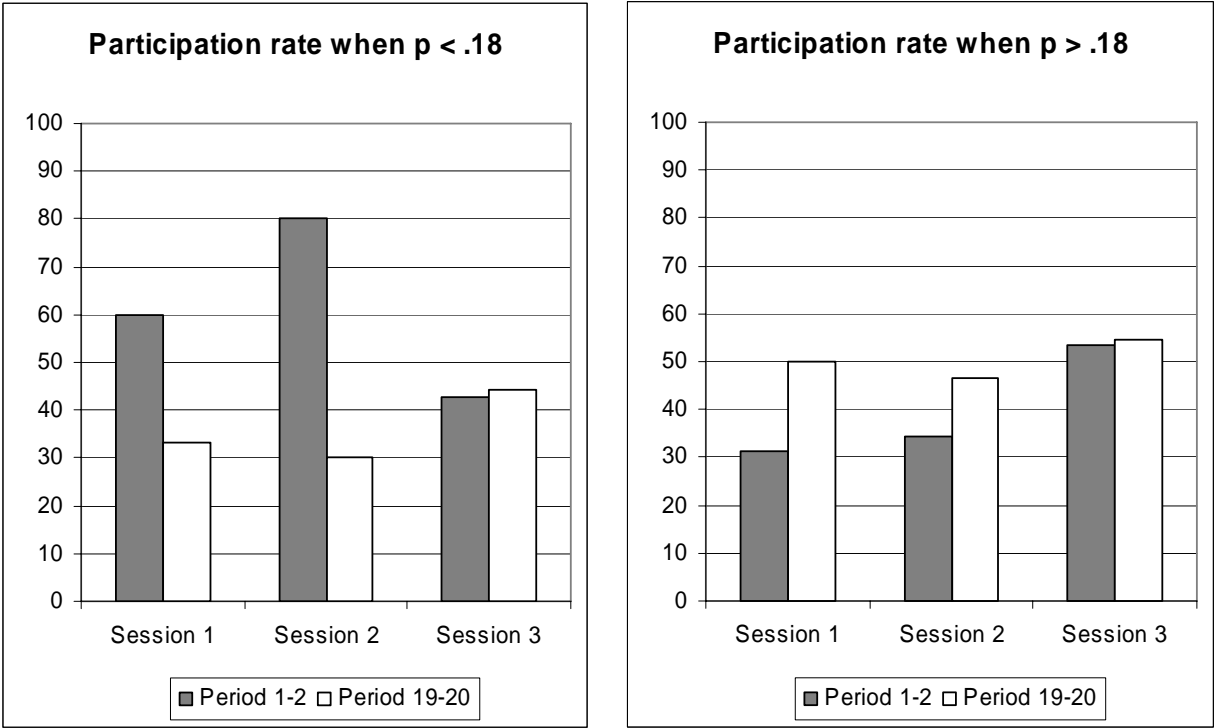


Figure 6: Average beliefs of being decisive (pivotal) over time for participants and non-participants

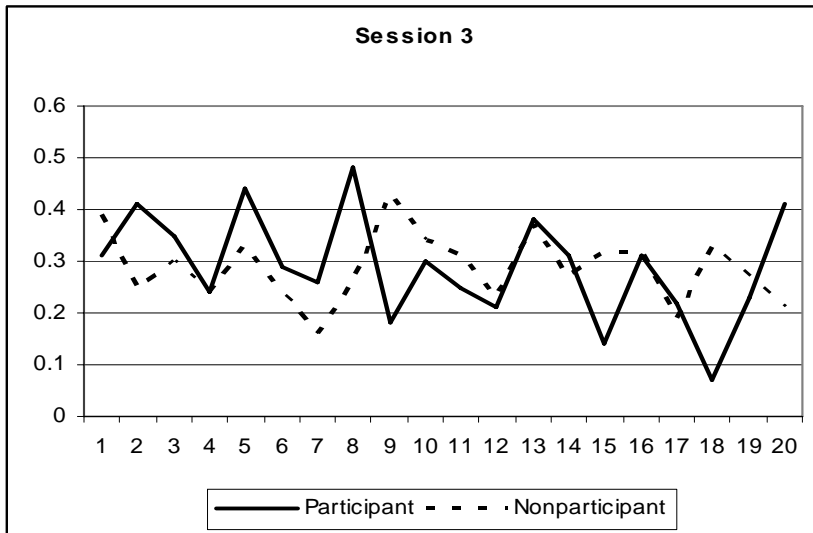




Figure 7: Turnout in advantaged and disadvantaged group

