

Research Article

Altruism, Noise, and the Paradox of Voter Turnout: An Experimental Study

Sarah A. Tulman

United States Department of Agriculture, Economic Research Service, 1400 Independence Avenue SW, Mail Stop 1800, Washington, DC 20250-0002, USA

Correspondence should be addressed to Sarah A. Tulman; sarah.tulman@ers.usda.gov

Received 4 December 2014; Revised 17 March 2015; Accepted 20 March 2015

Academic Editor: Arthur Schram

Copyright © 2015 Sarah A. Tulman. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

This paper addresses the paradox of voter turnout, wherein observed voting participation rates are far greater than what rational choice theory would predict. Voters face multiple voting choices, stochastic voting costs, and candidates offering different economic platforms. A combination of two approaches attempts to resolve this paradox: quantal response equilibrium (QRE) analysis, which introduces noise into the decision-making process, and the possibility of ethical (altruism-motivated) voting. A series of laboratory experiments empirically tests the predictions of the resulting model. Participants in the experiments are also given opportunities for communicating online with their immediate neighbors, in order to enhance the chances that subjects would realize the possibility of ethical voting. The results show that ethical voting occurs but gains momentum only in the presence of a vocal advocate and even then it mostly dissipated by the second half of the session. The QRE-based model was able to explain some but not all of the overvoting that was observed, relative to the Nash equilibrium prediction. There is evidence to suggest that communication via the chat feature generated some of the voting and also some of the ethical voting.

1. Introduction

The paradox of voter turnout, wherein observed turnout rates in elections are far greater than what is predicted under rational choice theory, has been a focus of researchers in economics and political science for over 50 years (see Geys [1] for a survey of the theoretical literature) and remains unresolved. Under rational choice theory, the probability of casting a pivotal vote—making or breaking a tie—is a key component of the motivation to vote. If the electorate is large then the probability of casting the pivotal vote falls to almost zero, and rational choice theory predicts that no one will vote. This probability declines quickly, so that even elections with relatively small electorates face this issue. And yet, observed turnout rates in elections are far greater than zero. For example, in the 2012 general election in the United States the estimated turnout rate among eligible voters was 58.6%. (All turnout rates are from the United States Election Project, <http://www.electproject.org>.) Even in years without a presidential election, when turnout is generally lower, it is still considerably higher than zero, such as the 41.8% turnout rate in the 2010 House and Senate elections. Therefore, there has

to be some other motivation for voting beyond the probability of one's vote being pivotal.

Why does this paradox matter? As the fields of election forecasting and polling analysis have expanded in order to take advantage of the explosion in available data, interest in these forecasts has rapidly increased. During the 2012 US presidential campaign there was an endless flow of political and polling analysis covered in almost every media outlet, and the presidential campaigns devoted vast resources to “get out the vote” efforts while using the rapidly increasing quantities of available data to microtarget voters. In spite of this increasing sophistication, one thing that can and often does cause election outcomes to diverge from forecasts is the matter of who actually votes on Election Day.

The model presented in this paper and the laboratory experiments which test its predictions focus on the economic motivations for voting. The payoffs that are offered by the candidates, and the equitability (or lack thereof) of these payoffs across groups, can be thought of as their economic platforms. Voters also face a cost associated with voting that models the explicit and implicit costs of going to the polls, such as transportation costs, lost wages due to taking time

off from work, or the opportunity cost of time. I draw from several threads of the literature which are at the intersection of economics and political science, and my model and experiments augment previous work on voting behavior.

The first thread from which my model and experiments are drawn focuses on the use of quantal response equilibrium (QRE), a generalization of the Nash equilibrium that allows for decision errors, to relax the assumption of perfect rationality. Building upon the framework presented in McKelvey and Palfrey [2], and also drawing from Goeree and Holt [3], Levine and Palfrey [4] used laboratory experiments to test voter turnout predictions in the presence of heterogeneous and privately known participation costs. However, in their paper the subjects could only choose between voting and abstaining, whereas in most elections potential voters can choose between at least two candidates, plus the option of abstaining. Would their results hold up if another choice were added so that potential voters could choose between three options: vote for one candidate, vote for the other candidate, and abstain? The model and experimental design presented here extend those of Levine and Palfrey [4] in this manner.

Another thread of the literature focuses on the role of ethical, or altruism-motivated, voting in generating voter turnout. Morton and Tyran [5] separated the motivations for voting choices into selfish voting and ethical voting, following the terminology used in the literature—voting for one's economic self-interest versus voting altruistically if one of the candidates offers a lower payoff but also a more equitable distribution of payoffs across groups of voters; however, their equilibrium predictions were qualitative in nature. By incorporating a similar payoff structure into the QRE-based framework from Levine and Palfrey [4], my model can generate quantitative predictions of turnout rates against which the experimental results can be compared. Additionally, it can be argued that incorporating bounded rationality and noise into the decision-making process may lead to a more realistic representation of a world in which individuals do not always (indeed, almost never) behave perfectly rationally. This combination was also used in Goeree et al. [6].

Another aspect of Morton and Tyran's [5] approach, albeit one that was necessary given the nature of the experiment, was that their experiment was a one-shot game. Although this is a more accurate representation of elections than a repeated game, in an actual election voters usually have several weeks or months leading up to the election during which voters can learn about and discuss the issues and the candidates' platforms. My experiments allowed for these learning effects during the experiment, through the inclusion of 20 rounds (elections).

My experiments also incorporated an online chat feature which allowed subjects to communicate before each round. The initial purpose of this was an attempt to increase the probability that the subjects would grasp the possibility of ethical voting, in treatments where that was relevant—the hope was that the subjects could clarify this (and other points) for each other. In the end, although this effect was indeed observed in certain cases, it also provided insight into subjects' motivations. The chat feature was also found to increase voting participation rates.

Group identity and communication within those groups have also been shown to increase voting participation, as seen in papers by Morton [7], Großer and Schram [8], Schram and van Winden [9], Schram and Sonnemans [10], and Charness et al. [11]. Kittel et al. [12], which was developed concurrently with this paper, incorporated multiple-candidate elections, costly voting as in Levine and Palfrey [4], and pre-voting communication. Although the focus and context of their paper differ from mine, they found that the distribution of earnings was more equitable when intergroup communication was allowed (analogous in some respects to the mixed-type chat groups in this paper). Their results also show that voting participation increases when communication is allowed, which matches the positive effect of communication on voting turnout that I found.

The main findings of this paper are summarized as follows. High rates of voter turnout were observed in the experimental data, especially when a subject voted for the candidate offering that subject's voter type a higher payoff than what the other candidate was offering. Some ethical voting was observed, but not enough to explain the overvoting. Also, most of the overvoting occurred when voters of each type voted in their own economic self-interest and, by its very nature, ethical voting is not a possible motive for that.

One feature that could potentially explain overvoting, unexpectedly, is the chat feature. (This feature was not included in the experimental design for the purpose of increasing voter participation but rather as a method to hopefully spread the word about the possibility of ethical voting and also to gain insight into the motivations behind decisions. This result was unexpected.) Because the chat was repeated before every round and the groups remained fixed throughout the session it set up a dynamic in which there may have been some accountability between chat partners. For example, one might ask the others whether they had voted in the previous round, and if so then for which candidate. The transcripts also showed that subjects would often urge others to vote in a certain way or would agree to vote in a certain way and sometimes would ask about whether/for whom others had voted in the previous round. Additionally, in spite of the differences in design and context, a few patterns were observed that had also appeared in the group identity literature, such as the emergence of subjects who persistently encouraged others to vote (analogous to Schram and van Winden's [9] producers of social pressure).

The rest of the paper proceeds as follows. Section 2 describes the theoretical model and results and the experimental design, which is influenced by the predictions generated by the theoretical model. Section 3 discusses the empirical results. Finally, Section 4 concludes.

2. Materials and Methods

2.1. Model. This model, in which voter turnout is a participation game (Palfrey and Rosenthal [13]), is an extension of the logit quantal response equilibrium (QRE) model used in Levine and Palfrey [4] and Goeree and Holt [3], incorporating the possibility of ethical voting along the lines of Morton and Tyran [5]. Whereas Levine and Palfrey [4] only allowed

TABLE 1: Payoffs (symmetric).

	Type 1	Type 2
Candidate X	q	$q - \theta$
Candidate Y	$q - \theta$	q

potential voters (also referred to simply as “voters,” unless a distinction is needed between actual and potential voters) to choose between voting and abstaining, here they have three choices: vote for Candidate X, vote for Candidate Y, and abstain. This better captures the reality of most elections, in which voters choose between two or more candidates in addition to the option of abstaining from voting. This extension of the model also allows for a range of payoff structures, including one which may induce ethical voting through the presence of a “selfish” candidate who offers a high payoff to one group of voters and a low payoff to the other and an “ethical” candidate who offers an equitable distribution of payoffs across groups of voters.

In this model, the voting rule is a simple plurality. Ties are broken fairly, with the winner being picked randomly (e.g., by a coin toss). Payoffs, the distribution of which can be thought of as the candidates’ economic platforms, are paid according to which candidate wins regardless of whether or for whom the participant voted. The payoff structure for each type is common knowledge to everyone regardless of type.

Potential voters are separated into two “types,” which determine the payoff that each one will receive depending on the winner of the election. Each candidate offers different payoffs to the two different voter types (Table 1). This table shows payoffs that are symmetric in both the payoff differences and the actual payoffs. This symmetry is not theoretically necessary—the main takeaway from this table is how the candidates offer different payments to each type—but matches the experimental design used in this research. With this payoff structure, voters of both types would face identical incentives to vote for the candidate offering their type the higher payoff.

Voting is modeled as being costly. This mirrors the costs associated with voting in real-world elections such as transportation costs, childcare, lost wages due to time off from work, or the opportunity cost of whatever someone would otherwise do. This cost is only incurred if someone votes. Costs are independent draws from the same uniform distribution, varying across both voters and elections, and both voter types face the same distribution of costs. In each election, each voter knows his or her own cost before deciding whether to vote but only knows the distribution from which others’ costs—and his or her own costs during future elections—are drawn.

Following Levine and Palfrey [4], a quasi-symmetric equilibrium in this model is a set of four turnout strategies $(\tau_{1X}, \tau_{1Y}, \tau_{2X}, \tau_{2Y})$ specifying the probabilities that a member of type i will vote for candidate j , as a function of the voting cost. For a quasi-symmetric equilibrium, it is assumed that within each type everyone will use the same strategy—although it is possible for members of each type to vote for either candidate, in the Nash equilibrium any voting will be

for the candidate that offers the higher payoff to his or her type.

The aggregate voting probabilities for each type-candidate combination are

$$p_{ij}^* = \int_{-\infty}^{\infty} \tau_{ij}(c) f(c) dc = \int_{-\infty}^{c_{ij}^*} f(c) dc = F(c_{ij}^*) \quad (1)$$

for each combination of voter type i and candidate j .

There is a cutpoint cost level c_{ij}^* below which, in the Nash equilibrium, the probability of voting equals one and above which the probability of voting is zero. At the cutpoint, the voter is indifferent. These cutpoints are expressed as

$$c_{ij}^* = \frac{(V_{ij} - V_{im})}{2} \times (\Pr(\text{make tie})_{ij} + \Pr(\text{break tie})_{ij}), \quad (2)$$

where V_{ij} is the payoff received by a member of voter type i if candidate j wins the election, and where m indexes the other candidate. This is the difference between the payoff received and the payoff that would have been received if the other candidate had won (the payoff difference). In the Nash equilibrium nobody would vote for the candidate offering one’s type the lower payoff because the payoff difference would be negative, implying a negative cost cutpoint in which case no cost would be low enough to induce voting. These cutpoints are set to zero because costs cannot be negative.

The term $(1/2) \times (\Pr(\text{make tie})_{ij} + \Pr(\text{break tie})_{ij})$ is the probability that a vote cast by a member of type i for candidate j will be pivotal (henceforth referred to as the “pivotal probability”). The one-half is there because, in the presence of two candidates, if someone’s vote *makes* a tie then there is a 50% chance that this voter’s chosen candidate will win (and that the vote will have been pivotal)—the winner of a tie is determined randomly. If someone’s vote *breaks* a tie then there is a 100% chance that his or her chosen candidate will win. However, because there was already a 50% chance that this candidate would have won in the event of a tie, there would only be a 50% chance that the vote that broke the tie changed that candidate’s outcome from losing to winning. Further details about the calculation of the pivotal probability can be found in Appendix A.

The expected payoff differences of voting for candidate j as a member of type i , net of costs, are expressed as

$$\pi_{k,ij} = (V_{ij} - V_{im}) \times \Pr(\text{pivotal})_{ij} - c_k, \quad (3)$$

where V_{ij} is the payoff received by a member of type i if candidate j wins the election and c_k is the voting cost faced by a given person in a given election. $\pi_{k,ij}$ can be thought of as a component of $\tilde{\pi}_{k,ij} = \pi_{k,ij} + \varepsilon_{k,ij}$, where $\varepsilon_{k,ij}$ is independent and identically distributed extreme value. This produces the logit model that is shown below.

Quantal response equilibrium (QRE) analysis, introduced in McKelvey and Palfrey [2], allows for “noise” in the decision-making process for voting. This process is often influenced by factors such as emotions, perception biases, voter error, unobserved individual heterogeneity, and other

sources of noise. The QRE introduces decision errors via a logit probabilistic choice rule according to which participants make their voting choices.

Under QRE, the best response functions are probabilistic rather than deterministic. Although “better” responses are more likely to be observed than “worse” responses, the “worse” responses are still observed because choice has a stochastic element. Thus, the assumption of perfect rationality is relaxed in favor of bounded rationality. As the amount of noise decreases—and the QRE approaches the Nash equilibrium—voters become increasingly likely to make choices that are consistent with the Nash equilibrium. In its most general form, the logit probability of choice “ x ” takes the form $P(x) = e^{\lambda\pi_x} / \sum_i e^{\lambda\pi_i}$, where π_i is the expected payoff from making choice i and λ is the logit precision parameter. The denominator $\sum_i e^{\lambda\pi_i}$ sums across all possible choices, summing to 1.

As the precision parameter λ increases, the level of noise decreases and the level of precision in decision making improves. As this happens, the ratio of “better” responses to “worse” responses improves. When λ is very large (very high precision/very low noise), the solution approaches the Nash equilibrium, with all or almost all decisions matching the Nash equilibrium outcome. (Strictly speaking, this convergence only holds for the principal branch of the QRE when there are multiple equilibria. However, for the model presented here the results were robust across a wide range of starting points.) As λ approaches zero (very low precision/very high noise), voting participation decisions approach randomness (50% if there are two choices, 33% if there are three choices, etc.) as the probabilities of “better” responses and “worse” responses become increasingly similar.

In the context of this model, a Nash equilibrium would correspond to a sharp drop in the response function at the cost cutpoint, where the probability of voting would equal one if facing a cost below c^* and would equal zero if facing a cost above c^* . For an intermediate level of noise, the cost cutpoint would affect the decision of whether or not to vote but the response function would not be as sharp because noise would lead to some instances of voting (or abstaining) even when the cost is above (or below) the cutpoint. A very low λ (very high noise level/very low level of precision) would correspond to a cutpoint that is effectively irrelevant, where the size of the cost would have very minimal or no effect on the decision of whether or not to vote. This can be seen in Figure 1.

The logit quantal response function, which can be thought of as the logit probability of a member of type i making voting choice j (where the choices are vote for X , vote for Y , and abstain) and facing cost k , is shown by

$$P_{k,\text{vote}ij} = \frac{e^{\lambda\pi_{k,ij}}}{e^{\lambda\pi_{k,iX}} + e^{\lambda\pi_{k,iY}} + e^{\lambda\pi_{k,ia}}}. \quad (4)$$

This can be simplified, for a type i voter choosing j (where the choices are only vote for X or vote for Y) and facing cost k , as

$$P_{k,ij} = \frac{e^{\lambda\pi_{k,ij}}}{1 + e^{\lambda\pi_{k,iX}} + e^{\lambda\pi_{k,iY}}}. \quad (5)$$

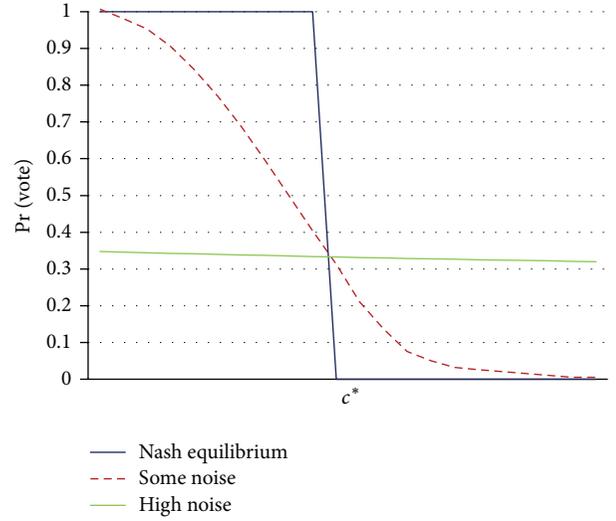


FIGURE 1: Logit response function relative to cost cutpoint c^* .

TABLE 2: Payoffs (asymmetric).

	Type 1	Type 2
Candidate X	$q + \theta$	$q - \theta$
Candidate Y	q	q

In the presence of random costs, these response functions have to be averaged across all possible costs (for a given type-candidate combination) in order to arrive at the average response probability of type i voting for candidate j . After averaging across costs, the probabilities are $(p_{iX}, p_{iY}, 1 - p_{iX} - p_{iY})$, where the third probability is the probability of abstaining.

The logit precision parameter, λ , is estimated using the likelihood function:

$$\begin{aligned} \log L = & \sum_l (N_{1X,l} \times \ln(p_{1X,l}) + N_{1Y,l} \\ & \times \ln(p_{1Y,l}) + (\text{num_obs}_{1,l} - N_{1X,l} - N_{1Y,l}) \\ & \times \ln(1 - p_{1X,l} - p_{1Y,l}) + N_{2X,l} \\ & \times \ln(p_{2X,l}) + N_{2Y,l} \times \ln(p_{2Y,l}) \\ & + (\text{num_obs}_{2,l} - N_{2X,l} - N_{2Y,l}) \\ & \times \ln(1 - p_{2X,l} - p_{2Y,l})), \end{aligned} \quad (6)$$

where $N_{ij,l}$ is the total number of votes by members of type i for candidate j that were observed during treatment l (or session) and $\text{num_obs}_{i,l}$ is the total number of observations for someone of type i in treatment l . These observations include decisions for any of the three possible choices: vote for Candidate X, vote for Candidate Y, and abstain.

One of the contributions of this paper is the incorporation of ethical (altruism-motivated) voting into a QRE framework. Table 2 shows the payoffs for an ethical candidate and a selfish candidate (following the terminology used in the literature).

The selfish candidate offers a high payoff to one voter type and a low payoff to the other type. The ethical candidate offers payoffs that have a more equitable distribution across types. In this respect, the payoff structure follows Morton and Tyran [5]. In Table 2, the ethical candidate offers payoffs that are equal across types. Also, the difference between the high and low payoffs is the same for both types. Neither of these features—the ethical candidate offering exactly equal payoffs or setting payoffs so that the payoff difference is the same for both types—are necessary, but it matches the experimental design that was used in this research.

In Table 2, q and θ stand for the baseline equitable payoff and the payoff difference, respectively. Candidate X is the selfish candidate and Candidate Y is the ethical candidate. Candidate Y 's payoffs maximize the minimum payoff and also minimize the difference between the payoffs received by the two types. Members of Type 1 receive a higher payoff if Candidate X wins but may be motivated by altruism to vote for Candidate Y even though they would receive a lower payoff if this candidate won the election. This is ethical voting. On the other hand, members of Type 2 receive a higher payoff if Candidate Y wins and would have no rational motive for voting for Candidate X because this candidate offers them a lower payoff *and* offers inequitable payoffs.

Selfish voting happens when economic self-interest is acted upon by voting for the candidate who offers one's type a higher payoff. This would happen when Type 1 votes for Candidate X or when Type 2 votes for Candidate Y . As discussed above, ethical voting happens when altruism outweighs economic self-interest and this would happen—at least potentially—when Type 1 votes for Candidate Y . It is also possible that this would occur because of noise. This is one example of how the chat feature proves useful for gaining insight into the motives behind voting decisions. There is no rational motive for Type 2 to vote for Candidate X ; therefore, this voting choice would be due to noise.

One difference between this design and Morton and Tyran's [5] is that in this design the aggregate payoffs are the same across candidates. I designed it this way so that, in the absence of ethical motives, equilibrium predictions would be the same across symmetric and asymmetric payoff structures for a given electorate size and a given noise level.

Ethical voting can be divided into two categories: ethical expressive and ethical instrumental. For ethical expressive voting, utility (or, in this model, implicit payoff) is gained from the act of voting for the ethical candidate regardless of the outcome of the election. This is along the lines of a warm glow parameter, as in Andreoni [14]. It is captured by the α term in the net expected payoff difference equation below for someone of Type 1 voting for Candidate Y and facing a person- and election-specific cost k :

$$\pi_{1Y} = (V_{1Y} - V_{1X}) \times \Pr(\text{pivotal})_{1Y} - c_k + \alpha. \quad (7)$$

Because ethical expressive voting is an additive term, it is independent of the pivotal probability. Therefore, it would increase the turnout rate even as the pivotal probability decreases, mitigating or even possibly outweighing the effect of the declining pivotal probability. The other three net

expected payoff difference equations, for the other type-candidate combinations, remain the same as in (3).

For ethical instrumental voting, utility is gained if the ethical candidate wins. Therefore, the utility/payoff gain is tied to the probability of one's vote for that candidate being pivotal. This is captured in the δ term in the equation below:

$$\pi_{1Y} = (V_{1Y} - V_{1X} + \delta) \times \Pr(\text{pivotal})_{1Y} - c_k. \quad (8)$$

In the presence of either ethical instrumental or ethical expressive voting, the estimation procedure for λ and α or for λ and δ is the same as before except that now two parameters are being estimated. In each case either α or δ , respectively, is identified through an exclusion restriction because the ethical voting parameter only appears in one of the four net expected payoff difference equations. Therefore, λ is pinned down by the other three net expected payoff difference equations and then α or δ is pinned down given λ .

In this model, it is not possible to simultaneously model both ethical expressive and ethical instrumental preferences. In that case, the net expected payoff difference equation would be

$$\pi_{1Y} = (V_{1Y} - V_{1X} + \delta) \times \Pr(\text{pivotal})_{1Y} - c_k + \alpha. \quad (9)$$

However, because $\alpha = \delta \times \Pr(\text{pivotal})_{1Y}$ there are an infinite number of possible combinations and α and δ cannot be separately identified.

2.2. Experimental Design. These theoretical predictions were empirically tested through a series of laboratory experiments. All sessions were conducted in the Veconlab Experimental Economics Laboratory at the University of Virginia on undergraduate students, using Veconlab software. There were 180 subjects—109 male and 71 female—across a total of 12 sessions. Each subject could only participate in one session; therefore, all subjects were inexperienced at the start of their respective sessions. Additionally, each session utilized only one treatment, so each subject only participated in one treatment. The average payout per subject was \$35, including a \$6 payment for showing up, and individual payouts depended upon individual decisions as well as the collective voting outcome.

Each session consisted of 20 rounds, or elections. Each subject was randomly assigned to one of two voting types, Type 1 and Type 2, at the start of the session and remained at that same type throughout the entire session in order to avoid reciprocity effects. Within each session, equal numbers of subjects were assigned to both types.

There were three possible choices: vote for Candidate X , vote for Candidate Y , and abstain. Anyone from either voter type could vote for either candidate. In the experiment, elections were referred to as rounds and candidates were referred to as options, as in "Option X " and "Option Y ," in an attempt to keep the language as neutral as possible. Throughout this paper, "round" and "election" will be used interchangeably and the options will be referred to as "Candidate X " and "Candidate Y " except in instances in which the word choice is relevant to the discussion at hand.

As discussed in the previous section, whichever candidate garnered the most votes within a round was declared the winner, with ties decided with a virtual coin flip. After the conclusion of each round, all subjects were told which candidate had won and the total number of votes for each candidate (but not the breakdown of how many of those were cast by members of each type). Each candidate offered different payoffs to each voter type, and subjects were paid the amount that had been promised to their type by the candidate who won in that round, regardless of whether or for which candidate they had voted. At the end of the session, subjects were paid their cumulative earnings.

In each round, each subject drew a voting cost from the same distribution. Costs ranged from \$0 to 0.42, in increments of \$0.02. This increment was chosen so that there were 22 possible costs and 20 rounds, meaning that, on average, there would be a good chance of a subject drawing almost all of the possible costs over the course of the session.

The upper bound of the range of costs was chosen to generate at least some abstention. Along the lines of Levine and Palfrey [4], the payoff from making or breaking a tie would equal $(V_{ij} - V_{im})/2$ for a member of type i voting for candidate j , where m is the other candidate. For both the asymmetric and symmetric treatments, the payoff from making or breaking a tie would equal \$0.25 for both types, according to the payoffs detailed below. Therefore, for subjects who drew a cost greater than this, abstention was the strictly dominant strategy. If there were a cost equal to \$0.25 then abstaining would be the weakly dominant strategy of a subject who drew this cost, but this was not a possible draw because of the \$0.02 increment size. The highest possible cost was set above the payoff from making or breaking a tie so that there would not exist an equilibrium in which the probability of voting was equal to one (i.e., in order to generate at least some abstention in equilibrium) Note that V_{ij} is assumed to be the higher payoff. As detailed earlier, there would be no cross-voting in the Nash equilibrium because the payoff differences would be negative, leading to a negative cost cutpoint.

At the end of the session, subjects participated in an exercise to elicit the subjects' perceived probabilities that their vote was/would have been pivotal in the final round, to test the relationship between voting turnout and the (perceived) probability of casting a pivotal vote. An analysis of the elicited probabilities showed a significant and positive relationship between the elicited probabilities and voting decisions—subjects who thought that they had a higher probability of being pivotal were more likely to have voted in the final round and with lower elicited probabilities were less likely to have voted. This section is available from the author upon request.

The subjects' computer screens showed short "backstories" next to each candidate's offered payoffs in the type-candidate payoff chart. The purpose of these was to help the subjects in the treatments with asymmetric payoffs realize the possibility of ethical voting. However, in order to avoid inadvertently encouraging or leading some subjects to vote in a certain way—as had happened in some of the pilot experiments—neutral wording was used here. The new backstories merely repeated the payoffs in words.

The backstories for asymmetric payoffs read as follows: "Option X will implement an investment that pays \$2 to Type 1 voters and \$1 to Type 2 voters" and "Option Y will implement an investment that pays \$1.50 to both Type 1 and Type 2 voters." For symmetric payoffs, the backstories were as follows: "Option X will implement an investment that pays \$1.50 to Type 1 voters and \$1 to Type 2 voters" and "Option Y will implement an investment that pays \$1 to Type 1 voters and \$1.50 to Type 2 voters."

Before the voting began in each round, subjects were able to chat online with their "neighbors" according to the ID number that had been assigned upon first logging into the program. For example, ID 3 could chat with ID 2 and ID 4. Meanwhile, ID 4 could chat with ID 3 and ID 5. The subjects with ID numbers at either end—ID 1 and ID 12 or ID 18, depending on the session size—were in the same chat group, creating a circular network. Subjects could also see the voter type of each of the other two subjects in their chat groups. The chat period lasted exactly one minute in each round, and it was not possible for anyone to submit a voting decision until after the chat ended.

This design models some features of real world discussions about politics. For example, most people only discuss an upcoming election with people whom they know directly (friends, family members, possibly coworkers, etc.). Even for those who broadcast their views on social media, the number of people who are on the receiving end of that is still small relative to the total size of the electorate. Given the small electorate size in this experiment, limiting a subject's chat circle to only two other people roughly mirrors this. Additionally, the overlapping group structure mirrors the fact that social groups are not self-contained, so that ideas or information can potentially be spread to a large number of people.

Furthermore, most of the chat groups contained members of both voting types. This was done to reflect the fact that most people have family members, friends, colleagues, acquaintances, and so forth who hold different political views than they do. (This differs from the focus of much of the group identity literature on the effects of communication within groups of voters.) Some subjects were part of a majority within their chat group and others were the minority, and in a few cases the subject was in a group consisting entirely of that subject's type.

Finally, the chat transcripts are extremely useful for providing insight into the subjects' motivations for voting a certain way or for abstaining altogether.

The experiment was divided into four treatments, which were designed to test the effects of electorate size—and thereby the role of the probability of casting a pivotal vote—and the possibility of ethical voting. First, the number of subjects per session—which can also be thought of as the size of the electorate—was varied. Half (six) of the sessions had 12 subjects per session, and the remaining sessions had 18 subjects. According to the underlying theory, the turnout rate should fall as the electorate size increases because the probability of one's vote being pivotal decreases. Second, the payoffs were varied so that in half of the sessions the subjects faced symmetric payoffs and in the other half the payoffs were

TABLE 3: Symmetric and asymmetric payoffs.

	Asymmetric		Symmetric	
	Type 1	Type 2	Type 1	Type 2
Candidate X	\$2.00	\$1.00	\$1.50	\$1.00
Candidate Y	\$1.50	\$1.50	\$1.00	\$1.50

asymmetric. These payoffs, which were common knowledge within a session, can be seen in Table 3.

In the treatments with asymmetric payoffs, Candidate X was the selfish candidate, offering a high payoff (\$2.00) to Type 1 subjects and a low payoff (\$1.00) to Type 2 subjects. The ethical candidate, Candidate Y, offered equal payoffs to both voter types. These payoffs maximize the minimum payoff and minimize the difference between the payoffs received by the two groups—indeed, in this design members of both groups would receive the same payoff. The payoff of \$1.50 for each group that was received by subjects of both voter types if Candidate Y won was lower than what Type 1 subjects would receive and higher than what Type 2 subjects would receive if Candidate X won. This design was inspired by Morton and Tyran [5].

To account for the possibility that Candidate X being listed first on the table (or X coming before Y alphabetically) might have affected voting decisions, the roles of ethical and selfish candidates were reversed in roughly every other asymmetric session so that Candidate X offered the equal payoffs and Candidate Y offered the inequitable payoffs. Everything else remained the same. No meaningful difference was observed in the data. In the final data, X and Y were reversed back for these sessions to make the data consistent with the other sessions.

Under this payoff design, members of Type 1 may face conflicting incentives in choosing a candidate. Type 1 subjects would maximize their earnings if Candidate X won but might also decide to vote for Candidate Y for altruistic reasons.

In the symmetric payoff design, the two candidates offered payoffs that were mirror images of each other. This meant that members of both voter types had the same economic incentives to vote for the candidate who was offering one’s type the higher payoff. Each candidate was equally inequitable, so there was no possibility of ethical voting here. These treatments served as a control, to measure the effect of noise in the decision-making process in the absence of any potential for ethical voting.

Each of the four treatments was a combination of session size (12 or 18) and payoff structure (asymmetric or symmetric).

The payoffs shown in Table 3 were set such that the equilibrium predictions of voting turnout rates were the same across asymmetric and symmetric payoff structures because the difference between the two candidates’ offered payoffs is identical across the two payoff structures so that the net expected payoff difference in functions from (3) remains the same.

Meanwhile, the number of subjects per session would affect the turnout predictions through the effect of the probability of one’s vote being pivotal on the net payoff function.

TABLE 4: (a) Predicted turnout probabilities and cost cutpoints (12 subjects). (b) Predicted turnout probabilities and cost cutpoints (18 subjects).

(a)			
	$\lambda = 100$ (Nash)	$\lambda = 7$ (L and P, 2007)	$\lambda = 0.2$ (high noise)
p_{1X}	0.2996	0.3069	0.3333
p_{1Y}	0	0.0793	0.324
p_{2X}	0	0.0793	0.324
p_{2Y}	0.2996	0.3069	0.3333
c_{1X}	0.122		
c_{1Y}	0		
c_{2X}	0		
c_{2Y}	0.122		

(b)			
	$\lambda = 100$ (Nash)	$\lambda = 7$ (L and P, 2007)	$\lambda = 0.2$ (high noise)
p_{1X}	0.2608	0.3069	0.3333
p_{1Y}	0	0.0793	0.324
p_{2X}	0	0.0793	0.324
p_{2Y}	0.2608	0.3069	0.3333
c_{1X}	0.122		
c_{1Y}	0		
c_{2X}	0		
c_{2Y}	0.122		

These both decrease (increase) as the electorate size increases (decreases). Therefore, the predicted turnout probabilities for a session with 18 subjects—regardless of whether it is a session with asymmetric or symmetric payoffs—are lower than those for sessions with 12 subjects, all else equal.

Tables 4(a) and 4(b) show the predicted turnout probabilities for the Nash equilibrium (extremely high λ), for $\lambda = 7$ (the value that was estimated in Levine and Palfrey, 2007), and for extremely high noise (very low λ). The predictions for a session with 12 subjects are shown in (a), and the predictions for a session with 18 subjects are shown in (b).

In Tables 4(a) and 4(b), p_{ij} is the predicted probability that a member of type i would vote for candidate j . The probability of a member of type i abstaining is $(1 - p_{iX} - p_{iY})$.

As shown in Tables 4(a) and 4(b), the probability of voting is increasing in the amount of noise (decreasing in the amount of precision). This is true for both own-voting and cross-voting, where own-voting is defined as voting for the candidate offering one’s type the higher payoff and cross-voting is defined as voting for the candidate offering one’s type the lower payoff. However, it is much more pronounced for cross-voting. In the Nash equilibrium, there is no cross-voting because in the presence of extremely high levels of precision (low noise levels) in the decision-making process, and in the absence of altruism, nobody would choose to vote for the candidate offering them a lower payoff because it would imply a negative expected payoff difference and a negative cost cutpoint. As the level of noise increases/the level of precision falls, it becomes increasingly likely that at least some of the subjects would cross-vote in spite of the negative expected payoff difference.

TABLE 5: Predicted turnout probabilities for different levels of ethical preferences (α).

For $\lambda = 7$	$\alpha = 0$	$\alpha = 0.1$	$\alpha = 0.25$	$\alpha = 0.5$	$\alpha = 0.75$	$\alpha = 1$
Type 1 for Candidate X	0.3071	0.2708	0.1606	0.0295	0.0052	0.0009
Type 1 for Candidate Y	0.0793	0.1575	0.3988	0.8232	0.9637	0.9934
Type 2 for Candidate X	0.0793	0.0818	0.1059	0.1629	0.1702	0.1707
Type 2 for Candidate Y	0.3071	0.2921	0.2347	0.175	0.171	0.1708
Candidate X (% of all voters)	0.1932	0.1763	0.13325	0.0962	0.0877	0.0858056
Candidate Y (% of all voters)	0.1932	0.2248	0.31675	0.4991	0.56735	0.5821

For a session with 12 subjects (Table 4(a)) the cost cutpoint is \$0.122, which means that there are seven possible costs—between \$0.00 and \$0.12—that are low enough that someone would vote (for the candidate offering one’s type the higher payoff) with probability of 1. Seven out of 22 possible costs equal approximately 30%, which is in line with the turnout probabilities in the Nash equilibrium. For a session of 18 subjects (Table 4(b)), the cost cutpoint is \$0.1047, which means that there are six possible costs—between \$0.00 and \$0.10—that are low enough to induce someone to vote. Six out of 22 possible costs equal approximately 27%, which is in line with the turnout probabilities in the Nash equilibrium.

One other thing to note is that as the noise level increases to the point where subjects are deciding randomly, the highest possible predicted turnout rate will be 33% (33% of each type would choose Candidate X, 33% would choose Candidate Y, and 33% would abstain). This is much lower than the 50% probabilities that would be the case if they could only choose between voting and abstaining, such as in Levine and Palfrey [4]. This will affect the ability of QRE to explain large amounts of overvoting, as will be discussed in the results section.

Additionally, the model predicts that increasing levels of ethical preferences can increase the percentage of the electorate voting for the ethical candidate and decrease the percentage voting for the selfish candidate (Table 5). (As in Table 4, the first four rows of Table 5 measure the percentage of that voters’ type who vote for a given candidate. The last two rows, capturing the voting decisions across voter types, measure the percentage of *all* voters.) As the level of altruism increases, ethical voting would reduce the closeness of the race, which in turn would decrease the probability of casting a pivotal vote. This loss of pivotality would not matter for the ethical voters, whose increasingly strong altruistic motivation would outweigh this, but it would lead those who would have voted for the selfish candidate to abstain.

3. Results and Discussion

There are several themes that appear throughout the results. The first is that voter turnout far exceeded even the highest possible turnout predicted by QRE. This was largely concentrated in votes for the candidate offering the higher payoff to that subject’s voter type and as such is not due to ethical voting. This indicates a strong economic self-interest. In addition, contrary to the predictions of the model—which were that, in the absence of altruism, the number of subjects (size of the electorate) would affect turnout but the symmetry of the payoff structure would not—the number of subjects

had no significant effect and turnout was significantly higher in sessions with symmetric payoffs relative to those with asymmetric payoffs. (However, the lack of a size effect is not entirely surprising considering that the theoretical differences in economically self-interested voting rates were not very large, as seen in Tables 4(a) and 4(b).) Third, some ethical voting (where motives were confirmed through the chat transcripts) and potentially ethical voting were observed; however, it was mostly found only under certain conditions and not enough was observed to explain the overvoting.

3.1. Terminology. Before getting into the analysis, some terminology that is used throughout this section needs to be defined. “Turnout” refers to all voting, regardless of type or candidate. “Favored” refers to votes for the candidate who offers that subject’s type a higher payoff than what the other candidate offers: Type 1 for Candidate X and Type 2 for Candidate Y. Similarly, “unfavored” refers to votes against their economic self-interest—Type 1 for Candidate Y and Type 2 for Candidate X. The unfavored category does not distinguish between potentially ethical voting—Type 1 voting for Y in a session with asymmetric payoffs—and cross-voting that is done for other reasons.

3.2. Voting Turnout Rates

3.2.1. Overview. The overall turnout rates, averaged across all rounds and broken down by type-candidate combinations, can be seen in Table 6(a).

The turnout rates for the candidates offering each type a higher payoff—Candidate X if Type 1 and Candidate Y if Type 2—are much higher than even the highest turnout rates predicted under QRE for any type-candidate combination, which topped out at 33% even in the presence of the highest possible levels of noise. Possible reasons for this overvoting, even relative to QRE predictions, will be addressed later. The second observation is that potentially ethical voting—Type 1 voting for Candidate Y in treatments with asymmetric payoffs—is higher than nonaltruistic cross-voting in these same treatments, which is Type 2 for Candidate X. In the treatments with symmetric payoffs, where there is no possibility of ethical voting, this is not the case.

There are two other notable relationships in the data. The first is that the turnout rate is higher, for all type-candidate combinations, in the treatments with symmetric payoffs than in those with asymmetric payoffs. This contradicts the model’s prediction, which was that there would be no difference in turnout between the two payoff structures

TABLE 6: (a) Voting turnout rates (all rounds). (b) Voting turnout rates (rounds 11–20). (c) Voting turnout rates (rounds 1–10).

(a)					
	All treatments	Symmetric	Asymmetric	12	18
Type 1 for Candidate X	0.4911	0.5467	0.4356	0.4500	0.5185
Type 1 for Candidate Y	0.0639	0.0656	0.0622	0.0597	0.0667
Type 2 for Candidate X	0.0544	0.0733	0.0356	0.0653	0.0472
Type 2 for Candidate Y	0.5250	0.5533	0.4967	0.5417	0.5139
(b)					
	All treatments	Symmetric	Asymmetric	12	18
Type 1 for Candidate X	0.4622	0.5022	0.4222	0.3972	0.5056
Type 1 for Candidate Y	0.0344	0.0533	0.0156	0.0167	0.0463
Type 2 for Candidate X	0.0211	0.0200	0.0222	0.0222	0.0204
Type 2 for Candidate Y	0.4911	0.5533	0.4289	0.5250	0.4685
(c)					
	All treatments	Symmetric	Asymmetric	12	18
Type 1 for Candidate X	0.5200	0.5911	0.4489	0.5028	0.5315
Type 1 for Candidate Y	0.0933	0.0778	0.1089	0.1028	0.0870
Type 2 for Candidate X	0.0878	0.1267	0.0489	0.1083	0.0741
Type 2 for Candidate Y	0.5589	0.5533	0.5644	0.5583	0.5593

except for possibly higher turnout for Type 1 voting for Candidate in treatments with asymmetric payoffs (potentially ethical voting). Also, there is no systematic difference in turnout between the two session sizes. This contradicts the model's predictions, in which turnout would be higher in sessions of 12 than in sessions of 18.

The turnout rates in the second half of the data (rounds 11–20) can be seen in Table 6(b).

This focus on the second half of the sessions (largely) sidesteps the effect of the learning curve that the subjects experienced in the early rounds. Once again, the effects of size and symmetry run counter to the model's predictions. Turnout is higher in treatments with symmetric payoffs relative to those with asymmetric payoffs, and session size has no systematic effect.

In rounds 11–20 the turnout rate for potentially ethical voting—Type 1 voting for Candidate Y in treatments with asymmetric payoffs—is even lower than the nonaltruistic cross-voting done by Type 2 for Candidate X, again in the asymmetric treatments. This points to altruism dissipating and economic self-interest taking over fairly quickly. Given that any possible altruism died out by the second half, it would be useful to look at the first half of the data (rounds 1–10). This is shown in Table 6(c).

Here we see potential altruism, with a turnout probability of 0.1089 for Type 1 subjects voting for Candidate Y in asymmetric treatments versus 0.0489 for Type 2 subjects voting for Candidate X in those same treatments. Also, the size and symmetry relationships are the same as before, albeit not as strongly.

It is also helpful to see the trends over time, not just the averages. Figure 2 shows the turnout rate—overall turnout, rather than being broken down into the different type-candidate combinations—over the course of the experiment.

Within Figure 2, the top row shows the turnout rates for 12 subjects (a) and for 18 subjects (b), and the bottom row shows the turnout rates for asymmetric payoffs (c) and symmetric payoffs (d).

A few trends appear in these figures. First, in all treatments the turnout rate decreases as the rounds go on. Second, there is no systematic difference between the 12- and 18-subject turnout rates. This is contrary to the theoretical results, which showed higher turnout for smaller electorate sizes. And finally, there is a clear difference between the treatments with asymmetric versus symmetric payoffs. As before, turnout in treatments with symmetric payoffs is almost always higher. This, too, is contrary to the theoretical results, which were identical across payoff structures for a given electorate size.

3.2.2. Nonparametric Permutation Test. A nonparametric permutation test (also known as a randomization test) was used to determine whether the turnout rates are different across treatments. This nonparametric test was useful when looking at session averages because there are only 12 data points. The results from the permutation tests showed the same trends that were observed in the stylized facts: the number of subjects per session does not significantly affect the probability of voting (except for symmetric sessions across all rounds), and subjects in treatments with symmetric payoffs are significantly more likely to vote than are subjects in treatments with asymmetric payoffs.

Details about the test and the results can be found in the Appendix.

3.2.3. Econometric Analysis. In the logit analysis, the dependent variables are the different categories of voter turnout—(“turnout”), favored, and unfavored. In each of the following

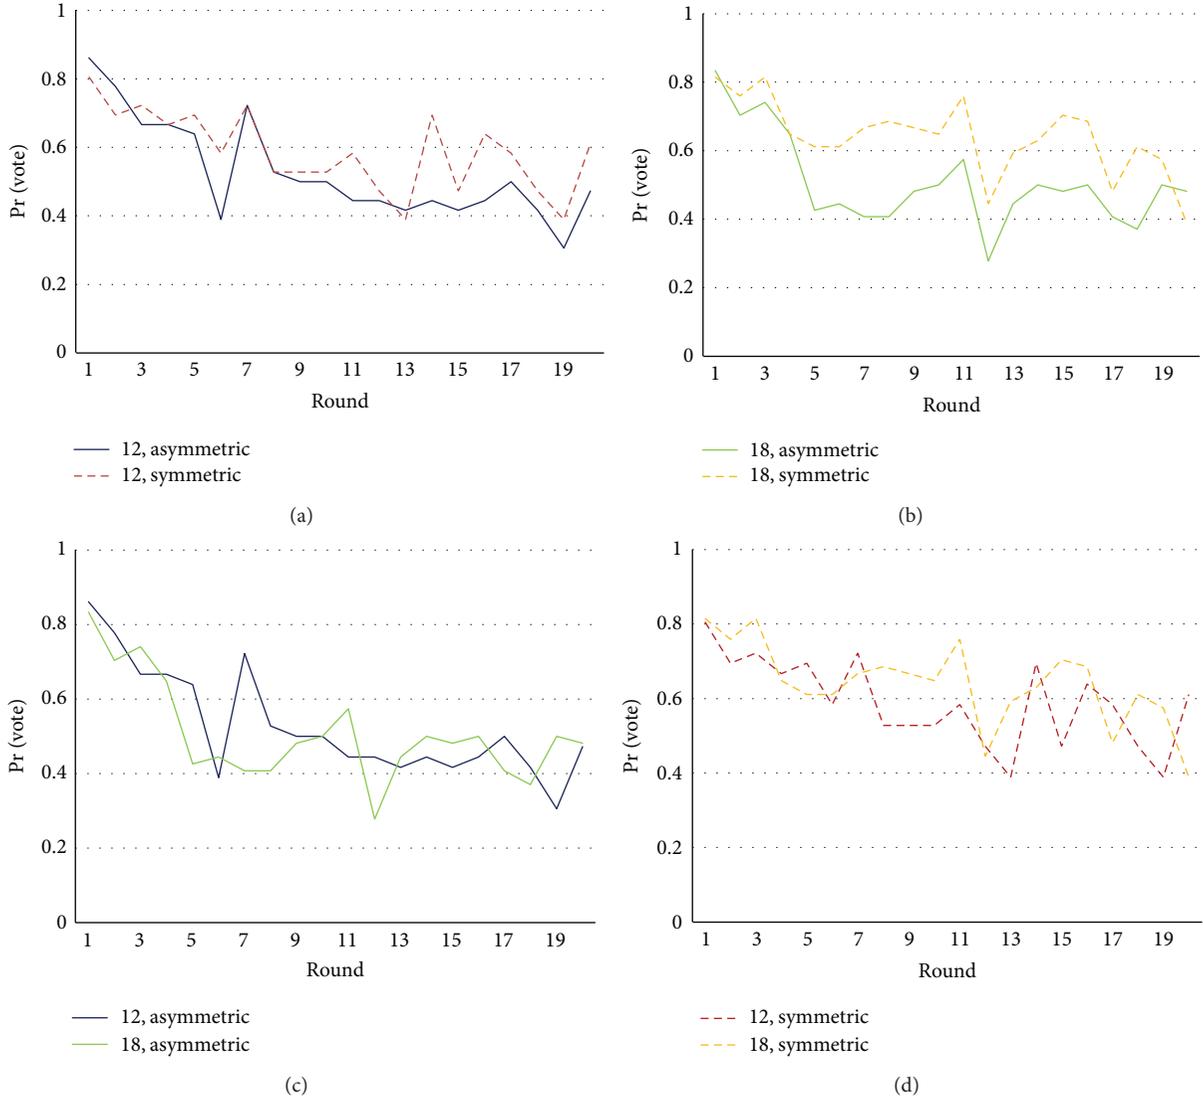


FIGURE 2: Voting turnout rates (top row symmetric versus asymmetric, bottom row 12 versus 18).

regressions, a random effects term was used to account for correlation within each subject's 20 decisions that cannot be attributed to any of the other independent variables—some subjects were inherently more (or less) likely to vote, or to vote in a certain way, than other subjects.

Table 7 shows the results for the following regressions:

$$\begin{aligned}
 \text{turnout}_{it} = & \beta_0 + \beta_1 (\text{average vote difference})_{it} \\
 & + \beta_2 (\# \text{ of previous losses})_{it} + \beta_3 \text{cost}_{it} \\
 & + \beta_4 \text{type}_{it} + \beta_5 \text{size}_{it} + \beta_6 \text{asymmetric}_{it} \\
 & + \beta_7 \text{round}_{it} + \beta_8 (\text{type 1} \times \text{asymmetric})_{it} \\
 & + \beta_9 (\# \text{ same type})_{it} \\
 & + \beta_{10} (\# \text{ group favored votes})_{it} \\
 & + \beta_{11} (\# \text{ group unfavored votes})_{it} \\
 & + w_{\text{turnout},i} + u_{\text{turnout},it}
 \end{aligned}$$

$$\begin{aligned}
 \text{favored}_{it} = & \gamma_0 + \gamma_1 (\text{average vote difference})_{it} \\
 & + \gamma_2 (\# \text{ of previous losses})_{it} + \gamma_3 \text{cost}_{it} \\
 & + \gamma_4 \text{type}_{it} + \gamma_5 \text{size}_{it} + \gamma_6 \text{asymmetric}_{it} \\
 & + \gamma_7 \text{round}_{it} + \gamma_8 (\text{type 1} \times \text{asymmetric})_{it} \\
 & + \gamma_9 (\# \text{ same type})_{it} \\
 & + \gamma_{10} (\# \text{ group favored votes})_{it} \\
 & + \gamma_{11} (\# \text{ group unfavored votes})_{it} \\
 & + w_{\text{favored},i} + u_{\text{favored},it}
 \end{aligned}$$

$$\begin{aligned}
 \text{unfavored}_{it} = & \theta_0 + \theta_1 (\text{average vote difference})_{it} \\
 & + \theta_2 (\# \text{ of previous losses})_{it} + \theta_3 \text{cost}_{it} \\
 & + \theta_4 \text{type}_{it} + \theta_5 \text{size}_{it} + \theta_6 \text{asymmetric}_{it}
 \end{aligned}$$

TABLE 7: Logit with subject-level random effects and cluster-robust standard errors (marginal effects).

	All rounds			Rounds 11–20			Rounds 1–10
	Turnout	Favored	Unfavored	Turnout	Favored	Unfavored	Unfavored
History: average $ X - Y $	-0.367*** (0.10)	-0.508*** (0.10)	0.037 (0.31)	-0.415* (0.23)	-0.435* (0.23)	-2.863* (1.36)	0.462 (0.34)
Number consecutive losses	-0.008** (0.00)	-0.012*** (0.00)	0.001*** (0.00)	-0.004 (0.00)	-0.009 (0.01)	0.000** (0.00)	-0.001 (0.00)
Type 1	-0.013 (0.06)	-0.002 (0.05)	-0.004 (0.01)	0.011 (0.06)	-0.017 (0.06)	0.000 (0.00)	-0.011 (0.01)
Cost	-1.554*** (0.13)	-1.620*** (0.12)	-0.343* (0.19)	-2.248*** (0.22)	-2.210*** (0.21)	-0.735** (0.36)	-0.254 (0.24)
Asymmetric	-0.129** (0.06)	-0.094* (0.06)	-0.020** (0.01)	-0.144** (0.06)	-0.140** (0.06)	-0.001 (0.00)	-0.031** (0.02)
Size 12	-0.022 (0.04)	-0.031 (0.03)	0.00 (0.01)	-0.04 (0.04)	-0.026 (0.05)	0.000 (0.00)	0.002 (0.01)
Round	-0.009*** (0.00)	-0.001* (0.00)	-0.002*** (0.00)	-0.005 (0.00)	-0.002 (0.00)	-0.000*** (0.00)	-0.004*** (0.00)
Type 1 * asymmetric	-0.023 (0.08)	-0.064 (0.08)	0.016 (0.01)	-0.034 (0.09)	-0.005 (0.09)	0.000 (0.00)	0.035* (0.02)
Number same type	-0.045 (0.04)	-0.032 (0.04)	-0.002 (0.01)	-0.001 (0.04)	0.025 (0.04)	-0.001 (0.00)	-0.002 (0.01)
Number favored	0.118*** (0.02)	0.163*** (0.02)	-0.013*** (0.01)	0.086*** (0.03)	0.095*** (0.03)	-0.000* (0.00)	-0.025*** (0.01)
Number unfavored	0.025 (0.02)	-0.019 (0.01)	0.005*** (0.00)	0.004 (0.02)	0.00 (0.02)	0.000 (0.00)	0.010** (0.01)
LR test (p value)	0.000	0.000	0.000	0.000	0.000	0.000	0.000

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Marginal effects elasticities used for continuous variables: history: average $|X - Y|$, cost. All others are marginal effects.

$$\begin{aligned}
 &+ \theta_7 \text{round}_{it} + \theta_8 (\text{type } 1 \times \text{asymmetric})_{it} \\
 &+ \theta_9 (\# \text{ same type})_{it} \\
 &+ \theta_{10} (\# \text{ group favored votes})_{it} \\
 &+ \theta_{11} (\# \text{ group unfavored votes})_{it} \\
 &+ w_{\text{unfavored},i} + u_{\text{unfavored},it}
 \end{aligned}
 \tag{10}$$

The last line in Table 7 is the likelihood-ratio test of whether the residual intraclass correlation equals zero (not a marginal effect, even though it was consolidated into this table). This tests whether the random effects term should be omitted. The results show that the use of random effects was appropriate.

Table 7 shows the marginal effects, where marginal effect elasticities were used for continuous variables and the usual marginal effects were used for binary or count variables. This is because “by what percentage does (dependent variable) increase if Asymmetric increases by 10%” is not meaningful, given that Asymmetric is binary. Note that the scales are different, so that an elasticity of 0.34 and a marginal effect of 0.034 (using an arbitrarily chosen example) would both be interpreted as a 0.34% increase in the dependent variable. This is why the results for “History: average $|X - Y|$ ” and “Cost” look so much larger than the others.

The independent variables are as follows. “History: average $|X - Y|$ ” measures the average, up until any given round, of the absolute value of the differences between the number of votes received by each candidate in previous rounds. This captures how close the outcomes of the elections had been and might have played a role in subjects’ perception of the probability of their vote being pivotal in the current round. The second variable, “number consecutive losses,” is the number of consecutive rounds, immediately preceding the round in question, that a subject’s type’s favored candidate (offering that type the higher payoff) has lost. This captures discouragement. “Type 1” is the voter type and equals one if Type 1 and zero if Type 2. “Cost” is the voting cost, between \$0 and \$0.42. “Asymmetric” equals one if the observation is from a session with asymmetric payoffs and zero if it is from a session with symmetric payoffs. “Size 12” equals one if the observation is from a session with 12 subjects and equals zero if it is from a session with 18 subjects. “Round” is the round number, between 1 and 20. “Type 1 * asymmetric” is an interaction term for a Type 1 subject in a session with asymmetric payoffs—the only combination for which ethical voting (Type 1 for Candidate Y) is possible—and equals one if the subject fits that description and zero otherwise. “Number same type” is the number of members of a subject’s chat group (excluding the subject) who are members of the same voter type as the subject. “Number favored” and

“number unfavored” are the number of members of a subject’s chat group (again, excluding the subject) who voted for the subject’s favored or unfavored candidate, respectively, in the current round. I also tested this using versions of “number favored” and “number unfavored” that use a rolling average of the previous 3 rounds, and the results were qualitatively similar in both magnitude and significance. Therefore, I decided to use the version above to focus on the more immediate effect of the most recent chat.

This analysis was done for all rounds and for the second half (rounds 11–20) for overall turnout, favored and unfavored, and for the first half (rounds 1–10) for unfavored only. It is uncommon to focus on the first half of the experimental data, because subjects often experience a learning curve during the first several rounds; however, in this case all of the ethical (motives confirmed through the chat transcripts) and potentially ethical voting occurred in the first half. This occurs in unfavored voting; therefore, this column is included in Table 7.

History: average $|X - Y|$, which measures the average number of votes separating the two candidates in previous rounds (the average margin), is significant and negative for both overall turnout and favored turnout, both for all rounds and for rounds 11–20. This means that the farther apart the vote counts had been in the previous rounds the less likely subjects were to vote, or conversely the closer the previous rounds had been the more likely subjects were to vote. Looking at the results for all rounds, if the vote counts in the previous elections were 10% farther apart the turnout rate would fall by 3.67% overall (or conversely, if previous elections were 10% closer the turnout rate would increase by 3.67%) and 5.08% for favored. The results in the second half were qualitatively similar, albeit less strongly significant possibly due to the average becoming less volatile as rounds go on. To make sure that this was not simply because the average incorporates more rounds as the session goes on, so that each incremental round has a smaller impact on the average, I also tried a rolling average of the previous three rounds and got qualitatively similar results.

The negative relationship between the size of the margin of victory in previous elections and the probability of voting in the current election (or the positive relationship between the closeness of previous elections and the probability of voting) may be because of the effect on subjects’ perceptions of the probability of their vote being pivotal. This result, and interpretation, qualitatively matches the results of the belief elicitation from the extended version of this paper (available upon request), which found that subjects’ *ex ante* beliefs about the probability of their vote being pivotal were positively related to the probability of voting. This, combined with the lack of size effect in the data, points to a possibility that voters’ *perceptions* of the probability could be the channel through which the probability of casting a pivotal vote impacts voting participation decisions, rather than the *actual* probability.

The number of consecutive losses immediately preceding a round is significant and negative for both overall turnout and favored turnout; the results show that with each additional consecutive loss the probability of voting falls by 0.9%

and the probability of voting for favored falls by 1.2%. For unfavored, the number of consecutive losses is significant but *positive*. It is not entirely clear why but may reflect the fact that as the number of previous losses adds up, the subjects who have a better understanding of the game may be more likely to abstain, leaving more voters who have a worse understanding and who would be more likely to make voting choices that go against their economic self-interest. This relationship weakens when focusing on only the second half because, in most sessions, losing/winning streaks were shorter in the second half.

The next significant variable is cost, which is highly significant and negative for all three dependent variables. For a 10% increase in voting cost, the probability of voting for either candidate falls by 15.54%, the probability of voting for favored falls by 16.20%, and the probability of voting for unfavored falls by 3.43%. The effect of cost on unfavored voting is both smaller in magnitude and in significance level, possibly because noneconomic factors would play a larger role in unfavored voting decisions and therefore the voting cost would have less impact. The cost effect is even stronger when focusing on only the second half, as subjects have gained a better sense of when it is or is not worth it to vote. In the first half, cost does not significantly affect unfavored voting, possibly because of the ethical (and potentially ethical) voting and the higher level of noise that occurred in the first half, both of which would have made subjects less sensitive to costs.

The results for size and symmetry match those from the permutation test and Figure 2, in that the probability of voting is lower in sessions with asymmetric payoffs relative to sessions with symmetric payoffs—for all rounds, subjects in sessions with asymmetric payoffs were 12.9% less likely to vote at all, 9.4% less likely to vote for favored, and 2.0% less likely to vote for unfavored—and the number of subjects per session is not significant. What could explain the higher turnout probability in sessions with symmetric payoffs, relative to those with asymmetric payoffs? There is no clear answer.

The results in Table 5 show that the model predicts increases in both unfavored voting and overall turnout as the level of altruism increases, but the relationship is negative in the empirical results. However, the drop in favored voting in Table 5 as altruism increases does match the direction of the empirical results. It is possible that, in the absence of any possibility for ethical voting, the subjects may have perceived the elections as being closer, which spurred greater amounts of voting. In the data, the elections were actually closer in sessions with asymmetric payoffs, but work with belief elicitation could potentially uncover a connection to the subjects’ *perceptions* of election closeness rather than the actual margins.

The interaction term of Type 1 and asymmetric captures the voters who have the option of voting ethically. The results here show that potentially ethical voting—Type 1 subjects voting for Candidate Y (unfavored) in sessions with asymmetric payoffs—is 3.5% more likely to occur than nonethical unfavored voting during the first half, which is when almost all of the ethical voting (both potentially ethical and transcript-confirmed ethically motivated) took

TABLE 8: Maximum likelihood estimates.

Parameters	Rounds	Sessions or treatments	Coefficient	Standard errors	<i>t</i> -statistics
λ	2nd half	All treatments	5.4909	0.5164	10.6332
	1st half	All treatments	1.0853	0.2841	3.8205
	1st half	Asymmetric treatments	1.732	0.4339	3.9704
	1st half	Symmetric treatments	0.5687	0.3767	1.5098
λ and α	1st half	Session 1, session 7			
		λ	0.5448	0.8978	0.6068
		α	0.129	0.3992	0.323
λ and δ	1st half	Session 1, session 7			
		λ	0.6155	0.8709	0.7068
		δ	0.9748	1.9373	0.5032

place. Note that all unfavored voting in sessions with symmetric payoffs and unfavored voting by Type 2 subjects in sessions with asymmetric payoffs are not even potentially ethical.

It is possible that ethical voting could have had a larger impact if the asymmetric payoff design featured higher aggregate payoffs. Fowler and Kam [15] found that altruists only had a larger incentive to participate, relative to those who are self-interested, when outcomes are perceived as benefitting everyone. Outcomes that are perceived as being merely distributive, without a larger aggregate benefit, are seen as being distributive and altruists gain nothing from merely shifting wealth from one group to another. In the experimental design used here, the aggregate payoffs are identical across candidates so the ethical candidate could be perceived as being distributive.

Round, which captures learning effects, is significant and negative for all three dependent variables when looking at all rounds and also for unfavored in rounds 1–10. As a session progresses the probability of voting falls by 0.9% (for overall turnout), 0.1% (for favored), or 0.2% (for unfavored) with each additional round. This points to the presence of a learning effect as the session proceeds, possibly as subjects figure out that they are less likely to be pivotal than they had anticipated or that it does not make sense to vote when facing a high cost. Then the size and the strength of the significance decline by rounds 11–20, indicating that the learning effect had dissipated because the steepest part of the learning curve happened early in the session.

The next three variables analyze the effect of subjects' chat groups on voting participation decisions. The group affiliation of the members of a subject's chat group did not significantly affect voting decisions, but whether/for whom their chat-mates voted did influence voting decisions. Looking at the results for all rounds, "number favored" increased favored voting by 16.3%, showing that economically self-interested voting participation was substantially increased by the influence of group members' voting decisions to vote for that same candidate. Unfavored voting *decreased* as the number of group members voting for a candidate's favored candidate increased, as voters who might have been leaning towards to vote against type were encouraged to vote for their favored candidate instead.

Unfavored voting *increased* in the number of group members who voted for that candidates' unfavored candidate, possibly due to factors such as group members encouraging each other to vote for the ethical candidate (where applicable), voting in solidarity with group members regardless of whether it is in the subject's best interests, and manipulation. Instances of these, as observed in the chat transcripts, can be found in Section 3.4. However, this 0.5% increase (or 1% in the first half) is much smaller than the effect of "number favored" on favored voting, possibly because it is an easier "sell" to convince people to vote for someone who is promising them more money than the other candidate.

3.3. Estimating the QRE and Ethical Voting Parameters. The model that was presented in Section 2 can be used to estimate the QRE precision parameter λ and the ethical voting parameters α and δ from the experimental data. To review, λ is a measure of the level of precision in the decision-making process of whether and for whom to vote. The ethical expressive parameter, α , measures the utility (or nonpecuniary expected payoff) received from the act of voting ethically. This is along the lines of a warm glow parameter, as in Andreoni [14]. The ethical instrumental parameter, δ , measures the utility (or nonpecuniary expected payoff) received from voting ethically if that vote is pivotal and leads to a victory for the ethical candidate.

Table 8 shows the results of the estimation using the experimental data.

First, λ was estimated for the second half (rounds 11–20). Ethical voting had dissipated by then, so limiting it to these rounds made it possible to estimate the QRE precision parameter in the (effective) absence of ethical voting. The estimated value for λ was 5.49, which is lower than what Levine and Palfrey [4] had estimated ($\lambda = 7$), meaning that there was more noise/less precision in the decision-making process here relative to Levine and Palfrey [4]. In that paper, subjects only had a choice between voting and abstaining so the higher noise/lower precision may be due to the addition of a third option (here, the possibility of choosing either candidate).

On a related note, it may also reflect the fact that voting turnout rates for Type 1 voting for Candidate X and Type 2 voting for Candidate Y are substantially above the highest

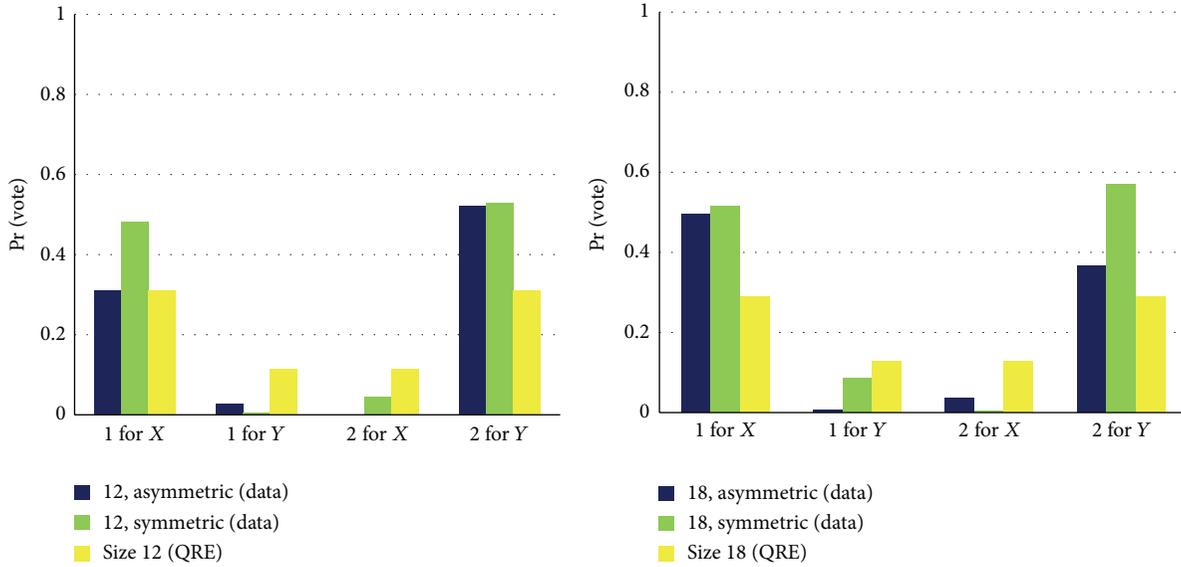


FIGURE 3: QRE versus data, rounds 11–20.

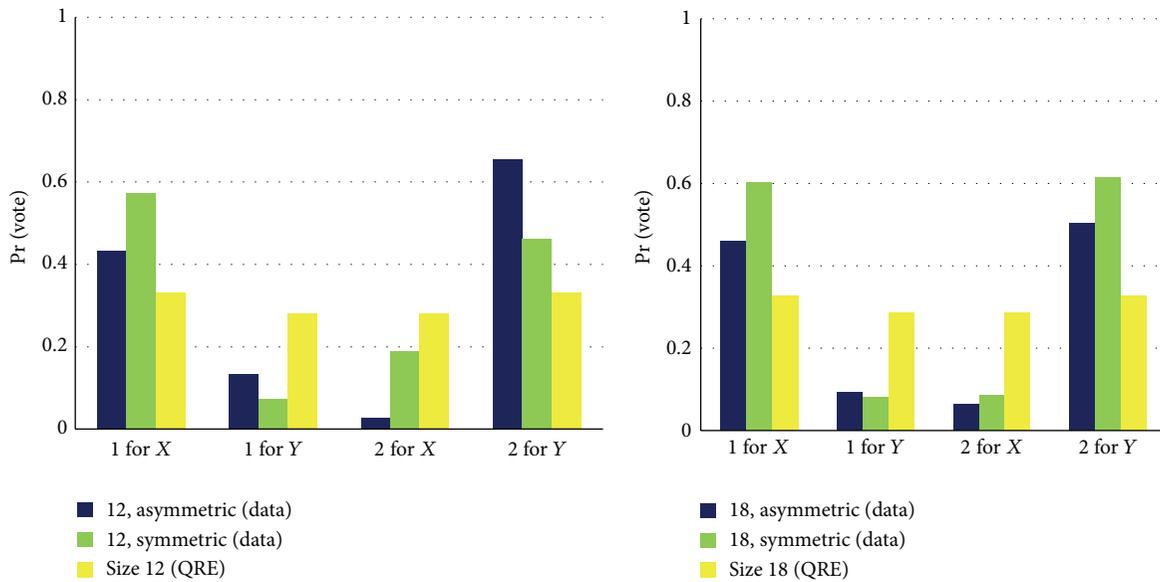


FIGURE 4: QRE versus data, rounds 1–10.

possible predicted turnout under QRE, which was 0.33 in the presence of extremely high noise as shown in Tables 4(a) and 4(b). This is much lower than the 0.5 predicted under high noise in Levine and Palfrey’s model (with only two options, when noise is so high/precision is so low that decisions are made randomly the probability of each option approaches 50%), leading to a smaller range within which overvoting could be explained by noise.

Figure 3 illustrates this overvoting, contrasting the QRE turnout predictions for 12 subjects and 18 subjects, given a λ of 5.49, with the observed turnout rates. For favored voting (Type 1 for X and Type 2 for Y), the observed turnout rates are much higher than the predicted rates, and the observed

turnout rates for unfavored voting (cross-voting) are lower than the predicted rates.

Next, as a comparison, λ was estimated for the first half of the data. As expected, the estimate of λ for the first half is considerably smaller—indicating lower levels of precision/higher levels of noise—than those in the second half. This is because of both higher turnout rates and greater levels of cross-voting. A comparison of observed turnout rates and QRE predictions is illustrated in Figure 4. Surprisingly, the estimated precision parameter is lower (higher noise) for sessions with symmetric payoffs, in spite of lacking any possibility of ethical voting. However, looking at the data, there are comparable levels of cross-voting, and the turnout rates of Type 1 for Candidate

X and Type 2 for Candidate Y are generally higher in the symmetric treatments. This higher turnout in favored voting is what is driving the difference. This is because the estimated λ is pulled lower by the high favored voting rates above and beyond what the model can explain, which is treated as extra noise.

However, the end goal is not only to estimate the QRE precision parameter λ but also to estimate the ethical voting parameters. This is done in the final two sections of Table 8. (As explained earlier, the ethical voting parameters α and δ cannot be separately identified so these are the results for separately estimated models, one with α and one with δ .) Unfortunately, for either of these specifications neither λ nor the respective ethical voting parameter is significant, even when only focusing on the two sessions with the highest rates of overvoting.

3.4. Chat. The online chat feature provides insight into the motivations behind subjects' voting decisions. (The chat transcript for all sessions is available upon request.) Two noteworthy themes appear in the chat: the first is the power of the overlapping chat circles to affect outcomes in both useful and harmful ways and the second is peer pressure/peer influence, even to the point of outright manipulation.

3.4.1. Transmission of Information. Although at least a small amount of potentially ethical voting was observed in most sessions with asymmetric payoffs, it appeared in a substantial fashion in Sessions 1 and 7. What differentiated these sessions from the other sessions with asymmetric payoffs was that in each case one of the subjects picked up on the possibility of ethical voting and voted accordingly and then broadcasted this realization and acted as a cheerleader of sorts to try to get Type 1 subjects to vote for Candidate Y . The "cheerleaders" were both Type 1 and Type 2, although their motivations differed according to type. These cheerleaders could be thought of as being thematically related to the producers of social pressure in Schram and van Winden [9]. In Session 1, ID 7—who was Type 1—started this chain of events in the chat before round 1. ID 7 picked up on the idea of ethical voting, "DO NOT BE GREEDY. WE WILL ALL MAKE MORE MONEY IF WE ALL CHOOSE THE LOWER OPTION." (All quotes appear exactly as in the chat transcript.) Strictly speaking, this was not entirely true because the aggregate payoff was the same regardless of which candidate won, and indeed Type 1 made *less* money when Candidate X won. However, the comment still demonstrates a desire to spread the payoffs more equitably. This subject then repeated this in the chat before round 2. Also in round 2, ID 8—who was a Type 2 and therefore was acting in his or her own economic self-interest—took up ID 7's argument by saying "Type 1 is being greedy," which spread the idea to ID 9—who was also a Type 2—who said, "Please choose option Y , we receive equal amounts regardless of type." Yet another Type 2, ID 11, said, "vote for option y . Otherwise Type 2 gets shafted." ID 10 was a Type 1 subject chatting with IDs 9 and 11 and started the chat period in round 2 by stating "Always vote for option X please. Same deal. We can get our money" but then changed his or her mind and ended up voting for Candidate Y after

being convinced by IDs 9 and 11. Of course, the fact that ID 10 said, "please" in the chat may indicate that this person was inherently more considerate and therefore maybe more likely to be open to the possibility of voting ethically, as opposed to someone like ID 5 whose sentiment was "sucks to be a 2."

Other rounds followed in a similar fashion. The amount of ethical voting never exceeded 2 subjects per round, and it was always done by IDs 7 and 10—and then only ID 7 after ID 10 started abstaining, but with only 12 subjects that was enough to make or break a tie in a few rounds. The chat also included some amusing attempts to generate ethical voting, such as "a vote for X is a vote for the bigwig corporations taking our hard-earned tax dollars," "a vote for X is un-American," and "here is a list of people who would vote for x in this scenario: bane, john lee hooker, ivan the terrible, judas."

In Session 7 a similar phenomenon occurred. (In Session 7, the selfish and ethical candidates had been switched so that Candidate X was the ethical candidate and Candidate Y was the selfish candidate. In this discussion the standard candidate name will be used, surrounded by brackets, rather than what was written in the transcript.) In round 3, ID 5 (Type 2) joked about the inequity of Candidate [X]'s offered payoffs. After this, ID 6 (Type 1) suggested voting for Candidate [Y] and ID 7 (Type 1) agreed. This block of voting continued for several more sessions. Also in round 3, ID 15 (Type 2) urged Type 1 subjects to vote for Candidate [Y] for ethical reasons, saying "but [y] makes money for everyone." (As explained earlier, this is not entirely true due to equal-sized aggregate payoffs but it demonstrates a desire for a more equitable distribution of payoffs.) However, it did not spread beyond this subject's original chat group.

One difference between Sessions 1 and 7 was the presence in Session 7 of voting that appeared to be ethical but was actually due to subjects not understanding how the payoffs and costs worked. This was the case for IDs 11 and 12, two Type 1 voters who banded together to vote for Candidate [Y] because they thought that it was to their advantage based on how high or low their costs were. Therefore, some of the seemingly ethical voting was actually due to noise or voter error. This is an example of how voting that appears ethical may not be and illustrates the usefulness of this chat feature for providing insight into motives.

It is also possible for incorrect information to spread and to affect the outcome of an election. In Session 12, ID 3—a Type 2 candidate—started off in the chat before round 1 by saying that everyone should vote for Candidate X . The chat period ended before he could follow up and explain why, but in round 2 the subject said, "continue with option X and forward it to your adjacent ID groups. If everyone votes X throughout, we all win." This strategy worked and by round 3 there were 11 votes for Candidate X , 1 vote for Candidate Y , and no abstentions.

The logic behind ID 3's strategy is a mystery because this was a session with symmetric payoffs so neither of the candidates was offering equitable payoffs. Also, as a member of Type 2 this subject received a *lower* payoff if Candidate X won. In round 5, this same subject wrote, "continue with x regardless of Type 1 or Type 2, it will all work out to

a guaranteed payout. fwd it to your group.” This subject had very quickly understood the power of the overlapping chat groups and yet apparently had not understood the instructions at the beginning of the session.

At the same time, word was spreading that subjects would not receive a payoff if they abstained or voted for the losing candidate. However, there were also individual conversations in which subjects were testing out the mechanics of the game. For example, in round 5 ID 8 had abstained and ID 9 had voted for Candidate Y, so in the chat before round 6 their discussion focused on whether payoffs were still received even if a subject abstained from voting. ID 8 wrote, “Yes. I did (abstain and still receive the payoff). It sucks that Option Y cannot win. But the better option now is to not vote if you’re Type 2.” ID 1 showed similar insight into this and also into the possibility of free riding in the presence of large margins for the winning candidate, “should we not vote this round? If the pattern stays the same, ppl will vote x and we won’t get charged.” Meanwhile ID 3 continued his or her push for Candidate X.

In the next round, there was chatter demonstrating continued confusion on the part of some subjects and among other subjects a growing understanding of the correct mechanics of the experiment. In this round five subjects ended up abstaining, indicating that the “always vote” strategy was breaking up. However, everyone who did vote chose Candidate X regardless of type. Evidently some were still confused about the link between payoffs and voting decisions, because in the chat before round 8 ID 5 said, “why would 5 people not vote and miss out on this cash cow?” ID 5 was a Type 1, so indeed it had been a cash cow for this subject except for the fact that he or she had incurred the voting cost in every round.

It was at this point that a rare event took place: the experiment was actually paused, and the rules and mechanics of payoffs and costs were explained again, such as the fact that the payoff received is a function of your type and the candidate that wins, not (directly) whether or for whom a subject voted. (This had already been stated several times during the instructions, which had been shown in writing and also read aloud at the start of the session. The instructions had also included a practice question that was specifically designed to make sure that subjects understood this and to correct any misunderstandings before the experiment began.) After this was explained again, voting decisions settled into a more normal pattern.

This incident provides an excellent example how the spread of information can affect outcomes, regardless of whether or not the information is correct. It is also an example of how relatively easy it is to exploit people’s uncertainty, as would have been the case at the beginning of the session, when there is always a learning curve regardless of how thorough the instructions are. It is important to note that this does not necessarily have to be done maliciously—it is unclear whether ID 3’s actions were due to misunderstanding or wanting to cause trouble.

3.4.2. Peer Pressure/Peer Influence/Manipulation. A second theme that was seen in the chat transcripts was that of subjects

urging each other to vote and sometimes agreeing to do so. This may have increased the turnout rates beyond what might have occurred without chat and may help to explain the high turnout rates that were observed. Also, because the chat is repeated over many rounds there is accountability to some extent—your chat neighbor might ask if you had voted and, if so, for whom you voted in the previous round. That sense of accountability could also increase voter turnout. And finally, some people feel a sense of solidarity in voting the same way as their friends/family/colleagues/and so forth.

In these experiments, this feeling of solidarity led subjects to vote against their own economic interests, even in the absence of altruistic motives. Usually this only lasted one or two rounds before self-interest took over. However, in a few cases this happened repeatedly. These present a possible parallel to Kittel et al.’s [12] swing voters who were influenced to vote differently than they would have otherwise. In this context, swing voter status was not due to experimental design but rather to certain subjects being more easily influenced than others.

One extreme case occurred in Session 10. ID 3, who was a Type 2 subject, was paired in a chat group with ID 2 and ID 4, both of whom belonged to Type 1. ID 3 wanted to vote with his or her chat group and voted for Candidate X—against the subject’s economic self-interest, but benefitting those of the other group members—in every round throughout the session. The subject asked whether they were all the same type before round 6, in spite of the types having been listed next to ID numbers in the chat interface in every round. However, even after ID 3’s group members clarified that they were both Type 2, ID 3 continued to vote for Candidate X. ID 3 voted for Candidate X in all but one round. Coincidentally, this extra vote for Candidate X either made or broke a tie in each of the first five rounds and set up a streak in which Candidate X won every round of the session, as subjects who would have voted for Y (mostly other Type 2s) got discouraged and stopped voting. I spoke with the subject after the end of the session to find out if this voting behavior resulted from a lack of understanding about the way in which the payoffs worked (being careful to ask about it in terms of wanting to find out whether the instructions were clear enough so that changes could be made, in an attempt to sidestep feelings of defensiveness that could have affected the answers). The subject had indeed understood everything but simply liked the feeling of solidarity from voting with the other chat group members.

In Session 11 there were two Type 1 subjects who often voted for Candidate Y. The payoffs in this session were symmetric, so altruism was not a possible motive. Both of them chatted with ID 15, a Type 2 candidate. ID 15 usually voted for Y but sometimes voted for X, usually when IDs 14 and 16 also voted X, and this subject spent much of the chat trying to convince his or her chat group to vote as a group. When asked about it after the session, IDs 14 and 16 each expressed a desire to vote with their chat groups, similar to the previous example. But then in a separate conversation ID 15 explained his strategy of occasionally voting for Candidate X—especially in the earlier rounds—for the purpose of generating a feeling of solidarity in IDs 14

and 16. This subject then would urge IDs 14 and 16 to vote for Candidate Y, which maximized ID 15's payoff. The end result was that ID 15 only voted for Candidate X four times, whereas ID 14 voted for Candidate Y eight times and ID 16 voted for Candidate Y twelve times, so on a net basis this strategy gained many more votes than it lost for Candidate Y.

This was the only incident of strategic manipulation that was observed in any of the sessions, at least as far as the transcripts indicated (and a bit of luck in running into the subject shortly after the session). However, it does illustrate the ability of a savvy political strategist to manipulate others who may not be as savvy, even those whose interests diverge from those of the strategist. It is also thematically related to the producers of social pressure in Schram and van Winden [9].

4. Conclusions

The analysis presented above attempted to resolve the paradox of voter turnout by incorporating ethical voting into a quantal response equilibrium- (QRE-) based model which allows for noise in the decision-making process and then testing the model's predictions using a series of laboratory experiments. The benefit of this modeling approach is that it generates predicted voting turnout probabilities, which were used in fine-tuning the experimental design, and against which the results of the experiments could be compared.

High rates of voter turnout were indeed observed. In fact, the turnout rates for some of the type-candidate combinations were so high that QRE-based analysis could not account for all of the overvoting, relative to the Nash equilibrium predictions. This is due at least in part to the presence of three voting choices—vote for Candidate X, vote for Candidate Y, and abstain—as opposed to the two choices of vote and abstain that were present in Levine and Palfrey [4]. Given a cost-payoff ratio that generates at least some abstention in equilibrium, the highest possible payoff occurs in the presence of extremely high levels of noise. With three choices, the probability of choosing any one of them is 0.33. This is considerably lower than the 0.50 when there are only two choices, which would not have explained all of the overvoting observed for some of the type-candidate combinations but would have at least explained more of it.

Ethical voting was not able to explain the overvoting, either. Some ethical voting was observed, but not enough to explain the extent of overvoting. Also, most of the overvoting occurred when subjects voted for the candidate offering their type a higher payoff. Ethical voting, which is defined as voting against one's own economic self-interest in order to vote for the candidate offering payoffs that are distributed more equitably across voter types, by its very nature does not fall into this category.

The online chat feature that subjects used for communicating with their virtual neighbors before each round (election) was shown to have increased voter participation. In this chat feature, subjects chatted online with the two subjects with adjacent ID numbers within the experiment. This mirrored the way in which people discuss politics

and upcoming elections with their friends, families, and coworkers.

Additionally, because the chat was repeated before every round it also set up a dynamic in which it may have generated some accountability between chat partners, for example, by asking one's chat group whether they had voted in the previous round, and if so then for which candidate(s). It was also observed in the chat transcripts that subjects would often ask others to vote in a certain way or agree to vote in a certain way and then sometimes follow up afterwards, which supports that conjecture. Through these channels, subjects' voting decisions were (mostly positively) influenced by the voting decisions of their group members.

There are several possibilities for future research in this area, such as comparing sessions with a chat feature versus no communication or the use of different chat group structures, for example, the current chat setup versus one in which subjects only communicated with members of their own voter type. Another possibility would be to focus on only ethical voting or only communication to clarify the individual impacts of these features on turnout rates and patterns. It is possible that the inclusion of both communication and (the potential for) ethical voting in this paper's experimental design obscured the importance of ethical voting. The chat feature was originally included for the insight that it would provide rather than being a tool to increase voting participation, and its impact on turnout was unexpected. Because of this, the experimental design did not carefully test the relative impacts of these two features. It would be interesting to design treatments to test for the individual effects, which I leave for future research.

Several other themes were also observed in the results, in which the patterns in the observed turnout rates differed from those that had been predicted. There was a strong cost effect on abstention rates, as predicted. However, there was no significant size effect, in contrast to the predictions in which turnout declined as the session size increased and the probability of casting a pivotal vote decreased. Second, the turnout rate was higher in sessions with symmetric payoffs (in which there was no possibility of ethical voting) than in sessions with asymmetric payoffs (in which there was a possibility of ethical voting). This contradicts the predicted outcome; however, the reason for this pattern is unclear. Belief elicitation might provide some insight into this but would require further research.

Appendices

A. Pivotal Probabilities

The term $(1/2) \times (\text{Pr}(\text{make tie})_{ij} + \text{Pr}(\text{break tie})_{ij})$ is the probability that a vote cast by a member of type i for candidate j will be pivotal (henceforth referred to as the "pivotal probability"). The one-half is there because, in the presence of two candidates, if someone's vote *makes* a tie then there is a 50% chance that this voter's chosen candidate will win (and that the vote will have been pivotal)—remember that the winner of a tie is determined randomly. If someone's vote *breaks* a tie then there is a 100% chance that his or her chosen

candidate will win. However, because there was already a 50% chance that this candidate would have won in the event of a tie, there would only be a 50% chance that the vote that broke the tie changed that candidate's outcome from losing to winning. Therefore, the sum of the probabilities of making a tie and of breaking a tie is multiplied by one-half.

The pivotal probabilities are analogous to those in Levine and Palfrey [4] and Goeree and Holt [3] but complicated by the fact that it is possible for votes received by each of the candidates to vote for either candidate to come from either type of voters. For example, even in the relatively straightforward example of 1 vote (total) for Candidate X and 1 vote (total) for Candidate Y, there are numerous possibilities for how this outcome arose: both votes came from Type 1 voters; Type 1 voted for Candidate X and Type 2 voted for Candidate Y; Type 2 voted for Candidate X and Type 1 voted for Candidate Y; and both votes came from Type 2 voters. Additionally, there is a separate pivotal probability for each type-candidate combination because it is possible for members of either type to vote for either candidate.

One example is the probability that a member of Type 1 will be pivotal by voting for Candidate X, modeled as a multinomial probability:

$$\begin{aligned}
 & \Pr(\text{pivotal})_{1X} \\
 &= \frac{1}{2} \left(\sum_{n_Y=1}^{\min(N_1, N_2)} \sum_{n_{2Y}=0}^{n_Y} \sum_{n_{2X}=0}^{n_Y-1} \binom{N_2}{n_{2X}, n_{2Y}} \right. \\
 & \quad \cdot \binom{N_1 - 1}{n_Y - 1 - n_{2X}, n_Y - n_{2Y}} \\
 & \quad \times p_{2X}^{n_{2X}} p_{2Y}^{n_{2Y}} (1 - p_{2X} - p_{2Y})^{N_2 - n_{2X} - n_{2Y}} \\
 & \quad \cdot p_{1X}^{n_Y - 1 - n_{2X}} p_{1Y}^{n_Y - n_{2Y}} \\
 & \quad \cdot (1 - p_{1X} - p_{1Y})^{N_1 - (n_Y - n_{2X}) - (n_Y - n_{2Y})} \\
 & \quad + \sum_{n_Y=0}^{\min(N_1, N_2) - 1} \sum_{n_{2Y}=0}^{n_Y} \sum_{n_{2X}=0}^{n_Y} \binom{N_2}{n_{2X}, n_{2Y}} \\
 & \quad \cdot \binom{N_2 - 1}{n_Y - n_{2X}, n_Y - n_{2Y}} \\
 & \quad \times p_{2X}^{n_{2X}} p_{2Y}^{n_{2Y}} (1 - p_{2X} - p_{2Y})^{N_2 - n_{2X} - n_{2Y}} \\
 & \quad \cdot p_{1X}^{n_Y - n_{2X}} p_{1Y}^{n_Y - n_{2Y}} \\
 & \quad \left. \cdot (1 - p_{1X} - p_{1Y})^{N_1 - 1 - (n_Y - n_{2X}) - (n_Y - n_{2Y})} \right), \tag{A.1}
 \end{aligned}$$

where N_i is the total number of potential voters of type i , n_j is the total number of votes for candidate j (cast by voters of either type), n_{ij} is the total number of votes by members of type i for candidate j , and p_{ij} is the probability of a member of type i voting for candidate j .

The first term within the brackets is the probability that a member of Type 1 will make a tie by voting for Candidate X. This means that before this incremental voting decision is made, there is one fewer vote for Candidate X than there is for Candidate Y, or $n_X = n_Y - 1$. The second term is the probability that a member of Type 1 will break a tie by voting for Candidate X. The fact that one vote would break a tie means that there is currently a tie and $n_X = n_Y$.

For each term within the brackets, the triple sum allows the calculation to sum over every possible combination of votes, for every possible number of votes. The outer sum loops over every possible value of n_Y (the total number of votes for Candidate Y), and therefore also over every possible value of n_X . The inner two sums loop over every possible combination of n_{2Y} and n_{2X} (the number of votes cast by members of Type 2 for Candidates Y and X, resp.). Because n_{1Y} and n_{1X} —the number of votes cast by members of Type 2 for Candidates X and Y, respectively—are simply $(n_Y - n_{2Y})$ and $(n_X - n_{2X})$, these summations also loop over every possible combination of n_{1Y} and n_{1X} .

In the first term, $\binom{N_2}{n_{2X}, n_{2Y}} p_{2X}^{n_{2X}} p_{2Y}^{n_{2Y}} (1 - p_{2X} - p_{2Y})^{N_2 - n_{2X} - n_{2Y}}$ is the probability that there will be exactly n_{2X} and n_{2Y} votes if all N_2 members of Type 2 have already made their voting (or abstention) decisions. $\binom{N_1 - 1}{n_Y - 1 - n_{2X}, n_Y - n_{2Y}} p_{1X}^{n_Y - 1 - n_{2X}} p_{1Y}^{n_Y - n_{2Y}} (1 - p_{1X} - p_{1Y})^{N_1 - (n_Y - n_{2X}) - (n_Y - n_{2Y})}$ is the probability that there will be exactly n_{1X} and n_{1Y} votes if all $N_1 - 1$ other members of Type 1 have already voted.

The second term follows the same format.

The other probabilities of a member of type i being pivotal by voting for candidate j equal

$$\begin{aligned}
 & \Pr(\text{pivotal})_{1Y} \\
 &= \frac{1}{2} \left(\sum_{n_X=1}^{\min(N_1, N_2)} \sum_{n_{2X}=0}^{n_X} \sum_{n_{2Y}=0}^{n_X-1} \binom{N_2}{n_{2X}, n_{2Y}} \right. \\
 & \quad \cdot \binom{N_1 - 1}{n_X - n_{2X}, n_X - 1 - n_{2Y}} \\
 & \quad \times p_{2X}^{n_{2X}} p_{2Y}^{n_{2Y}} (1 - p_{2X} - p_{2Y})^{N_2 - n_{2X} - n_{2Y}} \\
 & \quad \cdot p_{1X}^{n_X - n_{2X}} p_{1Y}^{n_X - 1 - n_{2Y}} \\
 & \quad \cdot (1 - p_{1X} - p_{1Y})^{N_1 - (n_X - n_{2X}) - (n_X - n_{2Y})} \\
 & \quad + \sum_{n_X=0}^{\min(N_1, N_2) - 1} \sum_{n_{2X}=0}^{n_X} \sum_{n_{2Y}=0}^{n_X} \binom{N_2}{n_{2X}, n_{2Y}} \\
 & \quad \cdot \binom{N_2 - 1}{n_X - n_{2X}, n_X - n_{2Y}} \\
 & \quad \times p_{2X}^{n_{2X}} p_{2Y}^{n_{2Y}} (1 - p_{2X} - p_{2Y})^{N_2 - n_{2X} - n_{2Y}} \\
 & \quad \cdot p_{1X}^{n_X - n_{2X}} p_{1Y}^{n_X - n_{2Y}} \\
 & \quad \left. \cdot (1 - p_{1X} - p_{1Y})^{N_1 - 1 - (n_X - n_{2X}) - (n_X - n_{2Y})} \right),
 \end{aligned}$$

$\Pr(\text{pivotal})_{2X}$

$$\begin{aligned}
&= \frac{1}{2} \left(\sum_{n_Y=1}^{\min(N_1, N_2)} \sum_{n_{1Y}=0}^{n_Y} \sum_{n_{1X}=0}^{n_Y-1} \binom{N_1}{n_{1X}, n_{1Y}} \right. \\
&\quad \cdot \binom{N_2-1}{n_Y-1-n_{1X}, n_Y-n_{1Y}} \\
&\quad \times p_{1X}^{n_{1X}} p_{1Y}^{n_{1Y}} (1-p_{1X}-p_{1Y})^{N_1-n_{1X}-n_{1Y}} \\
&\quad \cdot p_{2X}^{n_Y-1-n_{1X}} p_{2Y}^{n_Y-n_{1Y}} \\
&\quad \cdot (1-p_{2X}-p_{2Y})^{N_2-(n_Y-n_{1X})-(n_Y-n_{1Y})} \\
&\quad + \sum_{n_Y=0}^{\min(N_1, N_2)-1} \sum_{n_{1Y}=0}^{n_Y} \sum_{n_{1X}=0}^{n_Y} \binom{N_1}{n_{1X}, n_{1Y}} \\
&\quad \cdot \binom{N_2-1}{n_Y-n_{1X}, n_Y-n_{1Y}} \\
&\quad \times p_{1X}^{n_{1X}} p_{1Y}^{n_{1Y}} (1-p_{1X}-p_{1Y})^{N_1-n_{1X}-n_{1Y}} \\
&\quad \cdot p_{2X}^{n_Y-n_{1X}} p_{2Y}^{n_Y-n_{1Y}} \\
&\quad \left. \cdot (1-p_{2X}-p_{2Y})^{N_2-1-(n_Y-n_{1X})-(n_Y-n_{1Y})} \right),
\end{aligned}$$

 $\Pr(\text{pivotal})_{2Y}$

$$\begin{aligned}
&= \frac{1}{2} \left(\sum_{n_X=1}^{\min(N_1, N_2)} \sum_{n_{1X}=0}^{n_X} \sum_{n_{1Y}=0}^{n_X-1} \binom{N_1}{n_{1X}, n_{1Y}} \right. \\
&\quad \cdot \binom{N_2-1}{n_X-n_{1X}, n_X-1-n_{1Y}} \\
&\quad \times p_{1X}^{n_{1X}} p_{1Y}^{n_{1Y}} (1-p_{1X}-p_{1Y})^{N_1-n_{1X}-n_{1Y}} \\
&\quad \cdot p_{2X}^{n_X-n_{1X}} p_{2Y}^{n_X-1-n_{1Y}} \\
&\quad \cdot (1-p_{2X}-p_{2Y})^{N_2-(n_X-n_{1X})-(n_X-n_{1Y})} \\
&\quad + \sum_{n_X=0}^{\min(N_1, N_2)-1} \sum_{n_{1X}=0}^{n_X} \sum_{n_{1Y}=0}^{n_X} \binom{N_1}{n_{1X}, n_{1Y}} \\
&\quad \cdot \binom{N_2-1}{n_X-n_{1X}, n_X-n_{1Y}} \\
&\quad \times p_{1X}^{n_{1X}} p_{1Y}^{n_{1Y}} (1-p_{1X}-p_{1Y})^{N_1-n_{1X}-n_{1Y}} \\
&\quad \cdot p_{2X}^{n_X-n_{1X}} p_{2Y}^{n_X-n_{1Y}} \\
&\quad \left. \cdot (1-p_{2X}-p_{2Y})^{N_2-1-(n_X-n_{1X})-(n_X-n_{1Y})} \right). \tag{A.2}
\end{aligned}$$

TABLE 9: Voting turnout rates by gender.

	Male	Female
Turnout	0.54633	0.5993
Favored	0.50275	0.5162
Unfavored	0.04358	0.0831
Ethical	0.0375	0.01833

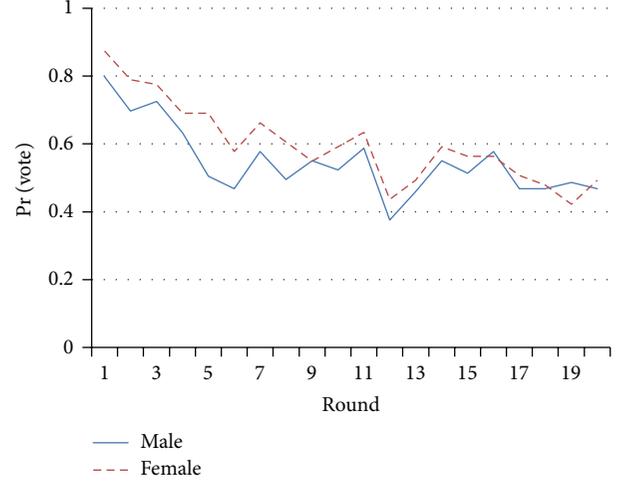


FIGURE 5: Voting turnout rates by gender.

These pivotal probabilities are then used to calculate the expected payoff differences, net of costs.

B. Gender Differences

Another area of interest is the effect of gender on voting decisions. In various nonparametric and parametric analyses gender was not found to have any significant effect on voting turnout, either for overall, for favored, or for unfavored. The turnout rates are summarized in Table 9 and Figure 5 illustrates the closeness of these turnout rates over time.

C. Nonparametric Permutation Test

Nonparametric permutation tests were used to test the significance of the difference between the means of two independent samples, when the sample sizes are small. The null hypothesis is that all of the observations are from the same population. Therefore, rejecting the null hypothesis would mean that all of the observations could not be from the same population, and the fact that the observations occurred specifically in their respective treatments cannot be attributed to coincidence. This was tested by reassigning the observations between the treatments, then comparing the differences in the means of the permutations with those in the original. The null hypothesis is rejected if the observed difference between treatments is large relative to the hypothetical differences in the other possible permutations.

First, looking at session-level average turnout rates for overall turnout, shown in Table 10, we see the same relationships as before in the averages for all rounds and for

TABLE 10: Voting turnout rates by session.

Treatment	Session	1st round	Average	Average (2nd half)
12, asymmetric	Session 1	0.8333	0.5625	0.4417
	Session 3	0.9167	0.5000	0.3750
	Session 9	0.8333	0.5208	0.4750
12, symmetric	Session 2	0.5833	0.6292	0.6500
	Session 4	0.8333	0.5292	0.4667
	Session 12	1.0000	0.6083	0.4750
18, asymmetric	Session 5	0.9444	0.5306	0.4778
	Session 7	0.7778	0.4778	0.4833
	Session 10	0.7778	0.5111	0.4000
18, symmetric	Session 6	0.7222	0.6306	0.6222
	Session 8	1.0000	0.6306	0.5556
	Session 11	0.7222	0.6583	0.5833

TABLE 11: Permutation test results (p value)—overall turnout rates.

	First round only		Average		Average (2nd half)	
	1-tailed	2-tailed	1-tailed	2-tailed	1-tailed	2-tailed
12: asymmetric versus symmetric	0.30	0.75	0.1*	0.200	0.150	0.300
18: asymmetric versus symmetric	0.50	1.00	0.05**	0.1*	0.05**	0.1*
Asymmetric: 12 versus 18	0.30	0.60	0.250	0.500	0.200	0.450
Symmetric: 12 versus 18	0.40	0.90	0.05**	0.1*	0.150	0.350
12 versus 18 (asymmetric and symmetric)	0.4142	0.8270	0.2767	0.5520	0.1357	0.2719
Asymmetric versus symmetric (12 and 18)	0.3419	0.6848	0.003***	0.006***	0.0038***	0.0078***

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

the second half (rounds 11–20). The first round of voting was also included because in most real-world elections one would only vote a single time.

The nonparametric permutation test then analyzes whether there really is a difference across treatments. The results are shown in Table 11. The results are reported as p values.

These results show the same trends that were observed in the stylized facts: the number of subjects per session does not significantly affect the probability of voting (except for symmetric sessions across all rounds), and subjects in treatments with symmetric payoffs are significantly more likely to vote than are subjects in treatments with asymmetric payoffs. Results for the first round only are not significant, which is not surprising. As opposed to an election outside of the lab, where potential voters have weeks or months to learn about the candidates and issues and to talk with their friends/family/colleagues, in the experiment there is quite a bit of learning that occurs during the first few rounds which may obscure any meaningful relationships in the data.

Next, the data for favored turnout and unfavored turnout were analyzed in the same manner. Given that most of the voting was for the favored candidate, the results found here were very similar to those for overall turnout. For unfavored turnout, no significant differences were found even after using the averages for rounds 1–10 (because that is when almost all of the ethical voting happened).

D. Instructions

The subjects saw the following set of instructions at the start of the experiment. These were also read aloud. The instructions below are from a 12-person session with asymmetric payoffs, so the instructions for the other treatments would vary with respect to the number of subjects and/or the payoff charts but would otherwise be the same. At the top of each page, the subject sees “Instructions (ID = —), Page — of 6”.

Page 1

- (i) *Matchings*. The experiment consists of a series of *rounds*. You will be matched with the *same* group of 11 other people in each round. The decisions that you and the 11 other people make will determine the amounts earned by each of you.
- (ii) *Voting Decisions*. At the beginning of each round, you will be asked to consider a vote for one of 2 options: Option X and Option Y.
- (iii) *Earnings*. The votes that are cast by the members of your group will determine a winning option, which will determine your earnings in a manner to be explained next.
- (iv) *Please Note*. Your earnings will depend on the outcome that receives the most votes (with ties decided at random), regardless of whether or not you voted for the winning option.

- (v) *Communications*. There will be a *chat room* open prior to the voting process, which allows you to send and receive messages. Voters will be identified by their ID numbers, and you are free to discuss any aspect of the voting process during the 1-minute chat period.
- (vi) *Communications Network*. You will be able to communicate with only two other voters in your group, who are your neighbors in the sense of having ID numbers adjacent to your ID. You should avoid inappropriate language or attempts to arrange side payments.

Page 2

- (i) *Voting Sequence*. The option that receives the most votes will determine earnings for all voters (irrespective of how each individual voted).
- (ii) *Ties*. In the event of a tie, one of the tied options will be selected at random, with each tied option having an equal chance of being selected.
- (iii) *Your Earnings*. The option selected in the voting process will determine your earnings. For example, your earnings for each outcome in the first round are

Option X: \$(*individual payoff*),
 Option Y: \$(*individual payoff*).

- (iv) *Voting Cost*. If you do vote, you will incur a cost. This cost will change randomly from round to round, and it will vary randomly from person to person. Each person's voting cost will be a number that is randomly selected from a range between \$0.00 and \$0.42. You will find out your voting cost before you decide whether to vote. You will not know anyone else's voting cost. Your total earnings will be the amount determined by the voting outcome, minus your voting cost (if any).

Page 3

- (i) *Individual Differences*. There are 12 voters in your group, who are divided into 2 "types," as shown in the table below. Your type and your earnings for each possible outcome are indicated in the *bright blue column*.
- (ii) *Payoff Backstory*.

Option X will implement an investment that pays \$2 to Type 1 voters and \$1 to Type 2 voters; Option Y will implement an investment that pays \$1.50 to both Type 1 and Type 2 voters.

Payoffs for all voters		
Outcome	type 1 (you)	type 2
Option X	\$2.00	\$1.00
Option Y	\$1.50	\$1.50
Number	6	
Voters:	(you and 5 others)	6

Page 4. Please select the best answer. Your answer will be checked for you on the page that follows.

Question 1. Suppose that your voting cost is \$0.21 and you are of *Type 1* in this round (different people will have different types). If you decide to vote and incur this cost, your earnings will be

- \$1.79 if Option X receives the most votes and \$1.29 if Option Y receives the most votes;
- \$1.79 if you voted for Option X and Option X wins, \$1.29 if you voted for Option Y and Option Y wins, and \$0.00 if the option you voted for does not win.

Submit answers

Payoffs for all voters

Outcome	type 1 (you)	type 2
Option X	\$2.00	\$1.00
Option Y	\$1.50	\$1.50
Number	6	
Voters:	(you and 5 others)	6

Page 5

Question 1. Suppose that your voting cost is \$0.21 and you are of *Type 1* in this round (different people will have different types). If you decide to vote and incur this cost, your earnings will be

- \$1.79 if Option X receives the most votes and \$1.29 if Option Y receives the most votes;
- \$1.79 if you voted for Option X and Option X wins, \$1.29 if you voted for Option Y and Option Y wins, and \$0.00 if the option you voted for does not win.

Your answer is *Correct*. You receive the payoff from the relevant (blue) column in the payoff table, regardless of whether you voted or not. In this case, the cost of voting is deducted since you voted in this example.

Continue

Payoffs for all voters

Outcome	type 1 (you)	type 2
Option X	\$2.00	\$1.00
Option Y	\$1.50	\$1.50
Number	6	
Voters:	(you and 5 others)	6

Page 6 (*instructions summary page*)

- (i) You will be matched with the *same* group of 11 other people in each round.
- (ii) All participants will begin by finding out their own voting cost, which will be a randomly determined number between \$0.00 and \$0.42. This cost will vary

from person to person. After seeing your voting cost, you will decide whether to incur this cost and vote. The option that receives the most votes will be selected. Your earnings will depend on the option selected, although your voting cost (if any) will be subtracted.

- (iii) *Please Note.* Your earnings will be determined by the outcome that receives the most votes (with ties decided at random), regardless of whether or not you voted for the winning option.
- (iv) Each round will begin with a 1-minute chat period. You will be able to communicate with only two other voters in your group, who are your “neighbors” in the sense of having ID numbers adjacent to your ID number. You will have to press the Update button manually to retrieve other’s messages as they arrive.
- (v) There will be a number of rounds in this part of the experiment. Your earnings for each round will be calculated for you and added to previous earnings.

Disclaimer

The views expressed are those of the author and should not be attributed to the U.S. Department of Agriculture or the Economic Research Service.

Conflict of Interests

The author declares that there is no conflict of interests regarding the publication of this paper.

Acknowledgments

This paper was based on a chapter from a Ph.D. dissertation completed at the University of Virginia under the helpful supervision of Charles Holt, Federico Ciliberto, and Simon Anderson. The author would also like to thank Craig Volden, Wayne-Roy Gayle, Beatrice Boulu-Reshef, Sean Sullivan, Angela Smith, Jura Liaukonyte, Jonathan Williams, Christopher Clapp, Nathaniel Higgins, Christopher Burns, James Williamson, Michael Harris, Charles Hallahan, and the participants of the Experimental Economics Workshop and the Political Psychology Working Group, both at the University of Virginia, for their helpful comments. The author is grateful to Arthur Schram and two anonymous referees for their comments and suggestions. In addition, the author thanks the Political Psychology Working Group, the Bankard Fund for Political Economy, and the National Science Foundation (Award no. BCS-0904795) for their financial support and the Department of Economics and the Veconlab Experimental Economics Laboratory at the University of Virginia for the experimental infrastructure.

References

- [1] B. Geys, “‘Rational’ theories of voter turnout: a review,” *Political Studies Review*, vol. 4, no. 1, pp. 16–35, 2006.

- [2] R. D. McKelvey and T. R. Palfrey, “Quantal response equilibria for normal form games,” *Games and Economic Behavior*, vol. 10, no. 1, pp. 6–38, 1995.
- [3] J. K. Goeree and C. A. Holt, “An explanation of anomalous behavior in models of political participation,” *American Political Science Review*, vol. 99, no. 2, pp. 201–213, 2005.
- [4] D. K. Levine and T. R. Palfrey, “The paradox of voter participation? A laboratory study,” *The American Political Science Review*, vol. 101, no. 1, pp. 143–158, 2007.
- [5] R. B. Morton and J.-R. Tyran, “Ethical versus selfish motivations and turnout in small and large electorates,” Working Paper, New York University, 2012.
- [6] J. K. Goeree, C. A. Holt, and S. K. Laury, “Private costs and public benefits: unraveling the effects of altruism and noisy behavior,” *Journal of Public Economics*, vol. 83, no. 2, pp. 255–276, 2002.
- [7] R. B. Morton, “Groups in rational turnout models,” *American Journal of Political Science*, vol. 35, no. 3, pp. 758–776, 1991.
- [8] J. Großer and A. Schram, “Neighborhood information exchange and voter participation: an experimental study,” *The American Political Science Review*, vol. 100, no. 2, pp. 235–248, 2006.
- [9] A. Schram and F. van Winden, “Why People Vote: free riding and the production and consumption of social pressure,” *Journal of Economic Psychology*, vol. 12, no. 4, pp. 575–620, 1991.
- [10] A. Schram and J. Sonnemans, “Why people vote: experimental evidence,” *Journal of Economic Psychology*, vol. 17, no. 4, pp. 417–442, 1996.
- [11] G. Charness, L. Rigotti, and A. Rustichini, “Individual behavior and group membership,” *American Economic Review*, vol. 97, no. 4, pp. 1340–1352, 2007.
- [12] B. Kittel, W. Luhan, and R. Morton, “Communication and voting in multi-party elections: an experimental study,” *The Economic Journal*, vol. 124, pp. 196–225, 2014.
- [13] T. R. Palfrey and H. Rosenthal, “A strategic calculus of voting,” *Public Choice*, vol. 41, no. 1, pp. 7–53, 1983.
- [14] J. Andreoni, “Impure altruism and donations to public goods: a theory of warm-glow giving,” *The Economic Journal*, vol. 100, no. 401, pp. 464–477, 1990.
- [15] J. H. Fowler and C. D. Kam, “Beyond the self: social identity, altruism, and political participation,” *The Journal of Politics*, vol. 69, no. 3, pp. 813–827, 2007.