Moral Bias in Large Elections *

Timothy Feddersen† Sean Gailmard‡ Alvaro Sandroni§

October 2006

Abstract

We provide theoretical and empirical support for the claim that large elections may exhibit a *moral bias*, i.e. alternatives that are understood by voters to be morally superior are likely to be chosen even when a majority of the eligible voters prefer another alternative. Using laboratory experiments we show that *ethical expressive* voters (voters who receive a payoff from taking an action they believe to be ethical) will have a disproportionate impact on election outcomes for two reasons. First, the choice of how to vote in a large election confronts voters with an essentially hypothetical choice — when ethical expressive types face hypothetical choice situations they are more likely to choose outcomes on the basis of ethical considerations than on the basis of narrow self-interest. Second, as pivot probabilities decline the set of people who participate will increasingly consist of ethical expressives.

*Paper prepared for the Annual Meeting of the Midwest Political Science Association, Chicago, IL, April 20-23, 2006. Thanks to Becky Morton for many helpful conversations and to Jenny Cheng for research assistance.

†Department of Managerial Economics and Decision Sciences, Kellogg School of Management, Northwestern University, 2001 Sheridan Road, Evanston, Illinois 60208.

‡Department of Political Science, Northwestern University, 601 University Place, Evanston, Illinois, 60208.

§Department of Economics, University of Pennsylvania, 3718 Locust Walk, Philadelphia PA 19104.

1 Introduction

A central argument for elections is that they produce results that are broadly representative of the preferences of a population. Consider a population that must choose between two options, A and B using an election. Suppose that each individual in the population would, if made dictator, choose B rather than A. In standard rational choice theory this is equivalent to saying that each person prefers B to A. In an election voters must decide whether to vote for A, B or abstain. In standard voting models (i.e., without asymmetric information) and no costs to vote each person has a weakly dominant strategy to vote for alternative B with the result that B would be the predicted election outcome. More generally, the standard theory predicts that election outcomes will be representative of the preferences of the population.[1] In this paper we provide a theory in which all voters may prefer B to A and yet vote for alternative A with the result that A wins the election. We provide evidence from laboratory experiments in support of our theory that suggests the predicted behavior is relevant for determining outcomes in real elections.

To illustrate the logic underlying our theory consider the case above and assume that B gives higher material benefit than A to each voter in an electorate. On the other hand assume everyone in the electorate agrees that A is morally superior to B. This might be the case if, for example, B gives high monetary returns to all voters while imposing high costs on a population of non-voters while A gives moderate benefits to voters and non-voters alike. Also assume that each voter if allowed to make a choice for the population would choose B. So, in spite of the moral superiority of A, by definition all voters prefer B to A.

In a standard model agents choose how to vote conditioning their vote choice on the event their vote is pivotal. Hence, such agents vote for the same alternative they would choose if they were dictator. Such is the case with *selfish* voters concerned only with maximizing their own monetary returns.

At the heart of our model is the assumption that agents get a positive subjective payoff for taking actions that they believe to be ethical. This payoff is received simply as a function of what action the agent chooses and

[1] Clearly, anything that may lead the preferences of the electorate to differ substantially from the preferences of the population as whole can lead to non-representative outcomes. For example, costs to vote are known to significantly decrease turnout (see Riker and Ordeshook 1968, Palfrey and Levine 2006) and may bias election results in favor of those with lower costs to vote. This paper provides an entirely different mechanism generating non-representative outcomes.

does not depend upon the outcome of the election. This is a significant departure from standard theory.[2] We call *ethical expressive* voters those who receive a payoff from taking an action they believe to be ethical.

Suppose that in the example above the electorate consists entirely of ethical expressive voters who get a small positive payoff from voting for alternative A. Now, if the payoff for acting ethically is small enough all voters prefer B to A. If the pivot probability is large enough then such ethical expressive voters will vote for B. However, in large elections the probability a vote is pivotal may be very small. When pivot probabilities are small, the voter's choice between A and B has very little impact on the actual outcome of the election. We say that such choices are *essentially hypothetical*.[3] In essentially hypothetical choice situations, the behaviorally relevant payoffs are those coming directly from the actions chosen. Ethical expressive voters, even though they prefer B to A, prefer to vote for A rather than vote for B.

This example shows that a very small payoff to act ethically can, in theory, have a large impact on election outcomes. The impact of ethical expressive types on election above would be small if, for example, a large fraction of the population consisted of selfish types. In that case all the selfish types would vote for B and B would be the outcome.

However, the impact of ethical expressive types is magnified when small costs to vote are introduced. It is well known that in large elections with costs to vote selfish types strictly prefer to abstain. Ethical expressive types will prefer to vote for A even in large elections with small costs to vote if the subjective payoff they get from voting for A is larger than the cost to vote. Thus, in large elections we may expect the electorate to consist almost entirely of ethical expressive types.

We test this theory by constructing an experiment in which a population chooses between one of two alternatives: A or B. The population is subdivided into A types who get a high payment if A is the outcome and nothing if B is the outcome. B types get a high payment if B is chosen and a smaller payment if A is chosen. A majority of the population are A types. The payments are chosen so that A maximizes the sum of payments and gives nearly equal payments to everyone. For these reasons we call A the *ethical* alternative.

One way to test our theory would be to set up an experiment in which

[2] See Feddersen and Sandroni (2006a and b) for a discussion of such models.

[3] A choice that has no impact on the actual outcome of the election may reasonably be called a *hypothetical choice.*

subjects vote in an election and pivot probabilities are varied implicitly by changing the size of the electorate and other parameters. The difficulty with such an approach is that in voting games there are typically multiple equilibria and so it is not clear how perceived pivot probabilities change as the size of the electorate increases. To better control pivot probabilities, we simulate an election with costly voting by a decision mechanism in which a subset of individuals are designated as *active* and may either choose to vote for A or B at some cost ($c > 0$) or abstain at no cost. The outcome is determined when one active individual is selected at random. If the selected individual has not abstained, his vote determines the outcome. If the selected individual has abstained then one of the two options is chosen by a flip of a fair coin. The number of active individuals determines precisely the probability that an active individual's vote choice is pivotal. If there is only one active individual then that person is a dictator. If there are n active individuals the probability an active's vote is pivotal is simply $1/n$.

In our experiment only B types are active. This is done because we want to focus on the impact of changes in pivot probabilities on the incentives for people to vote against their material interests. The parameters of the experiment are chosen so that active B types who care only about maximizing their own payments have an incentive to vote for B only if they are a dictator. When the pivot probability is less than .5 selfish types strictly prefer to abstain rather than incur the cost to vote.

Ethical expressive voters receive the same monetary payments as selfish types but are modeled as receiving a positive subjective payoff by voting for the ethical alternative A. Depending upon the magnitude of the subjective payoff ethical expressive types may have an incentive to vote for B when the pivot probability is high, to abstain for moderate pivot probabilities and to vote for A when pivot probabilities are low. In contrast to selfish types, as pivot probabilities decrease ethical expressive types may have increasing probabilities of voting and may switch their vote choice from B to A.

We also model an alternative ethical voter type we call an *ethical instrumental* voter. Such types are modeled as getting an additional positive payoff whenever the ethical alternative A is the outcome. Individuals motivated by altruism are an example of such types. While ethical instrumental types may vote for alternative A they, like selfish types, have an increasing incentive to abstain as pivot probabilities decrease and their vote choice is insensitive to pivot probabilities.

When voting costs are positive the main behavioral predictions for ethical expressive, ethical instrumental, and selfish voters are: (1) the incentive for ethical expressive types to vote for A is either constant or increasing as pivot probabilities decrease, while for the other two types the probability of voting for either A or B is decreasing as pivot probabilities decrease; and (2) all three voter types have decreasing incentives to vote for B as pivot probabilities decrease. When costs to vote are zero (3) the incentive to vote for A is not affected by changing pivot probabilities for either the selfish or instrumental types but a decrease in pivot probability results in an increased incentive for ethical expressive types to vote for A .

In a series of laboratory experiments we find empirical support for hypotheses (1) and (2) above and that ethical expressive types represent an economically relevant proportion of the electorate (i.e., their individual behavior significantly affects collective choices). We have not yet tested (3).

Our results suggest that large elections may exhibit a *moral bias* i.e., alternatives that are understood by voters to be morally superior are likely to be chosen even when a majority of the eligible voters may prefer another alternative.

A central contribution of this paper is to suggest that election outcomes may be morally biased because of two separate impacts of low pivot probabilities on the decision on whether and how to vote. First, the choice of how to vote in a large election confronts voters with an essentially hypothetical choice–when ethical expressive types face hypothetical choice situations they are more likely to choose outcomes on the basis of ethical considerations than on the basis of narrow self-interest. Second, as pivot probabilities decline the set of people who participate will increasingly consist of ethical expressives.

The paper is organized as follows. We first provide a brief review of the literature on turnout in elections. Then we first discuss the implications of several alternative theoretical models of behavior for the participation decision. Following this we introduce the experiment design and specific predictions of each model. Section 5 presents the findings from the experiment and discusses their implications for the theoretical models. Finally, Section 6 concludes.

2 Review of the Literature

Literature on Turnout and Civic Duty

Literature on Warm Glow Giving

Voting Experiments with Costly Voting

Literature on hypothetical versus real choices (to be done)

3 Basic Concepts and Theoretical Predictions

Before we describe the experimental design in detail it will be helpful to describe the game-theoretic substructure of the experiment. Consider a group N consisting of $n > 0$ individuals that must choose between two options, A and B. The group is composed of two subgroups, A types who get a higher monetary reward if option A is the outcome and B types who get a higher monetary reward when option B is the outcome. Let $n_A > 0$ and $n_B > 0$ denote the number of individuals of each type where $n_A + n_B = n$.

The set of B types is further subdivided into *active* and *inactive* individuals. Let n_β be the number of active B types and $n_{\tilde{}\beta}$ be the number of inactive B types so that $n_\beta + n_{\tilde{}\beta} = n_B$. Only active B types have a chance to influence the group decision.

Active B types simultaneously and privately choose one of three options: abstain, vote for A or vote for B. The group decision is determined by selecting one active B individual at random. If the selected individual has voted then his vote determines the outcome. If he has abstained then the group outcome is determined by a the flip of a fair coin.

Payoffs are given below:

Table1.Monetary Rewards under options A and B.

	A type	active B type who vote	other B types
Option A	$1 - c$	$1 - c$	1
Option B	0	$1 + x - c$	$1 + x$

The term c and x are parameters in the model where $c > 0$ corresponds to a monetary cost of voting for active B types. The parameter x corresponds to a monetary premium for B types if option B is the outcome. A types receive a monetary reward of $1 - c$ if alternative A wins the election and 0 otherwise.

We assume that $1/2 > x > 2c > 0$ and $n_A > n_B$. These assumptions ensure that alternative A minimizes inequality in terms of monetary rewards, maximizes the sum of monetary rewards, and gives a higher monetary reward to a majority of the group. For these reasons we say that A is the *ethical outcome*.[4]

[4] In fact the assumption $x > 2c$ could be replaced by the weaker assumption that $x > c$. However, this stronger condition simplifies the exposition and is consistent with the monetary payoffs we offered in our experiments. We discuss the case $2c > x > c$ in an appendix.

3.1 Selfish Types

We call types who only care about maximizing their own expected monetary rewards *selfish* types. For simplicity, we assume that the payoff of the selfish type is identical to the monetary reward that he receives. It is easy to see that this simple model produces incentives to vote that are entirely analogous to the incentives in standard models of elections. To see this we show that as the probability a vote is pivotal decreases the incentive to vote also decreases.

The payoff to the active B type for voting for B is

$$\frac{1}{n_\beta}(1+x)+\left(1-\frac{1}{n_\beta}\right)(1+q^*x)-c$$

where $\frac{1}{n_\beta}$ is the probability he is selected to be decisive (i.e., the probability that his vote is pivotal) and q^* is the probability option B is chosen when his vote is not pivotal. With probability $\frac{1}{n_\beta}$ the voter is pivotal. In that case because he voted for B he receives the payoff of $1+x-c$. With probability $\left(1-\frac{1}{n_\beta}\right)$ the voter is not pivotal. In that case he receives a payoff $1+q^*x-c$. Because voting is simultaneous q^* is independent of the voting decision made by the voter. The payoff to the active B type from abstaining is

$$\frac{1}{n_\beta}(1+\frac{x}{2})+\left(1-\frac{1}{n_\beta}\right)(1+q^*x)$$

Note that when the voter abstains he does not pay the cost of voting and when he is pivotal half the time B is the outcome. Thus, the selfish B type weakly prefers to vote for B rather than abstain if and only if

$$\frac{x}{2n_\beta}\geq c.$$

So, as the probability a vote is pivotal $(\frac{1}{n_\beta})$ decreases the incentive for a selfish B type to abstain gets larger.

For purposes of the empirical analysis below we note that conditional on choosing to vote, a selfish B type has a strictly dominant strategy to vote for B.

3.2 Ethical Instrumental Voters

Voters may depart from selfish behavior if they take into account the monetary rewards of others. A large body of experimental literature suggests that such considerations are important. We define *ethical instrumental* voters as those who prefer option A to option B.

We model ethical instrumental voters as receiving an additional payoff $\delta > x$ when the ethical option A is chosen. For an ethical instrumental voter the only payoff difference between voting for A and B occurs when his vote is pivotal. In that case he prefers to vote for A since $\delta > x$. As with the selfish type (who always prefers B) the ethical instrumental voter's choice between A and B is independent of the probability his vote is pivotal.

The determination of when the ethical instrumental type votes (as oppose to abstains) is entirely analogous to the analysis with selfish types. The payoff to this type for voting for option A is

$$\frac{1}{n_\beta}(1+\delta) + \left(1 - \frac{1}{n_\beta}\right)(1 + q^* x + (1-q^*)\delta) - c$$

while the payoff for abstaining is

$$\frac{1}{n_\beta}(1 + \frac{x}{2} + \frac{\delta}{2}) + \left(1 - \frac{1}{n_\beta}\right)(1 + q^* x + (1-q^*)\delta)$$

Hence, ethical instrumental voters will prefer to vote for A if and only if

$$\frac{\delta - x}{2n_\beta} \geq c.$$

So, as the probability of being pivotal decreases (n_β increases) the incentive for an ethical instrumental voter to participate decreases.

3.3 Ethical expressive voters

Feddersen and Sandroni (2005, 2006) develop a model in which ethical voters get a payoff for taking an action they determine to be ethical independent of the consequences of that action. We call agents who get a payoff simply by voting for option A whether or not their vote is pivotal *ethical expressive* types.

Ethical expressive types get the same payoffs as selfish voters plus a payoff of $d > c$ by voting for option A.[5] The payoff to this type for voting for option A is

$$\frac{1}{n_\beta} + \left(1 - \frac{1}{n_\beta}\right)(1 + q^* x) + d - c$$

[5] The assumption $d > c$ ensures that the behavior of the ethical expressive voters is different qualitatively from the behavior of the selfish voters. We relax this assumption in the appendix.

while the payoff for voting for option B is

$$\frac{1}{n_\beta}(1+x)+\left(1-\frac{1}{n_\beta}\right)(1+q^*x)-c$$

Conditional on voting, ethical expressive voters prefer to vote for A over B if

$$d\geq\frac{x}{n_\beta}.$$

So, conditional on voting, as the probability of being pivotal decreases the incentive for an ethical expressive type to vote for A increases.[6] Note that this is in contrast to both the selfish and ethical instrumental models where pivot probabilities does not impact the choice between A and B.

Voters with $d\geq\frac{x}{n_\beta}$ prefer to vote for A rather than abstain if and only if

$$\begin{aligned}&\frac{1}{n_\beta}(1)+\left(1-\frac{1}{n_\beta}\right)(1+q^*x)+d-c\\ \geq\;&\frac{1}{n_\beta}(1+\frac{x}{2})+\left(1-\frac{1}{n_\beta}\right)(1+q^*x)\end{aligned}$$

or

$$d-c\geq\frac{x}{2n_\beta}.$$

Voters with $d<\frac{x}{n_\beta}$ prefer to vote for B rather than abstain if and only if

$$\begin{aligned}&\frac{1}{n_\beta}(1+x)+\left(1-\frac{1}{n_\beta}\right)(1+q^*x)-c\\ \geq\;&\frac{1}{n_\beta}(1+\frac{x}{2})+\left(1-\frac{1}{n_\beta}\right)(1+q^*x)\end{aligned}$$

or

$$\frac{x}{2n_\beta}\geq c.$$

With some algebra it can be shown that the behavior of ethical expressive types is as follows. Consider three different cases (d large) $d\geq x$; (d intermediate) $x>d>2c$; and (d low) $2c>d$. When d is large the ethical expressive voter always votes for A. When d is intermediate this type votes for B when the pivot probability is high ($\frac{1}{n_\beta}>\frac{d}{x}$), votes for A otherwise.

[6] As the pivot probability increases the set of pairs (d,δ) that satisfy the equation above decreases (by inclusion).

When d is low then this type votes for B when the pivot probability is large ($\frac{1}{n_\beta} > \frac{2c}{x}$), abstains when the pivot probability is in the interval ($\frac{2(d-c)}{x}$, $\frac{2c}{x}$) and votes for A is small ($\frac{2(d-c)}{x} > \frac{1}{n_\beta}$).

Ethical expressive types behave much differently from either selfish or instrumental types. Such types may exhibit both a propensity to vote for the selfish alternative B when pivot probabilities are high and a propensity to vote for the ethical alternative A when pivot probabilities are low. This may seem counterintuitive but it has a straightforward intuition. As pivot probabilities decrease the choice of which candidate to vote for become essentially *hypothetical* in the sense doesn't have much impact on the voter's material payoff. Therefore the potential benefit from voting selfishly becomes small while the subjective payoff from voting for the ethical alternative stays constant.

A second behavioral difference between ethical expressive types and selfish or instrumental types is that in the former case the incentive to vote may be increasing as pivot probabilities decrease whereas in the latter cases the incentive to vote is decreasing. To understand this phenomena consider a voter who must choose between abtaining and voting for A. When pivot probabilities are high, abstention relative to voting for A produces a significant increase in the probability the selfish alternative B is chosen. However, when pivot probabilities are very low the choice to abstain does not significantly change the probability that B is chosen.

From this analysis one can predict the effects of pivot probability on participation by agents as a function of their motivations, and therefore the effects of pivot probability on the probability that each option is chosen for the group.

4 Experiment Design

The experiment design described below allows us to test the predictions from Section 3 in a controlled setting.

The experiment was conducted is a sequence of rounds. A round can be broken into a sequence of four stages. In stage 1, a group N of n subjects is partitioned into two subsets $N_A, N_B \subset N$ corresponding to A and B types. The size of each subset is n_A and n_B respectively. Further, a subset $N_\beta \subset N_B$ of size n_β of the B types are designated as *active types.* Each subject in a group is informed of the number of people of each type before any decisions are made. Subjects know which category they themselves are in but are not informed of the identity of other individuals in these categories. A B type learns whether he or she is an active type before making any decisions.

In stage 2, each active type must choose whether to vote or not. If he chooses to vote then he pays a small cost c and specifies one of the two outcomes A or B. All other subjects have no decision to make.

In stage 3, after all active types make their participation and vote choice, one active type is randomly selected from the set of all active types. The probability any active type is selected is $\frac{1}{n_\beta}$ which is the probability an active type is pivotal. Note that any active type can be randomly selected at this stage, whether they have chosen to vote or not.

Stage 4 determines the group choice. If the active type selected at stage 3 has chosen to vote, then the outcome that subject specified at stage 2 is the group choice. If this voter has not voted then the outcome, A or B, is chosen by a fair coin toss.

The sequence of four stages makes up a single round of a session of the experiment. After one round is completed then another begins with a new random draw of A, B and active types. A sequence of rounds with groups drawn from a set of participants comprises a session of the experiment. The treatment variables subject to experimental control are n_A, n_B, and n_β. As noted n_β determines the probability a vote is pivotal while changes in n_A and n_B determine the collective benefits that result from each outcome. A sequence of rounds with fixed values for n_A, n_B, and n_β in a session is a distinct *treatment* in the experiment.

Payoffs in the experiment are determined as in Table 1 above. In all rounds of the experiment, $c = .10$ denotes the participation cost and $x = .25$ denotes the premium that B types earn from option B over option A. Participants are informed of these parameters in the instruction period and in a table visible to them at all times in the experiment.[7]

The experimental design does not use a control group; instead results from all treatments are aggregated and we analyze the effect of each design variable on individual and group choices. In addition, both the subjects and the experimenters are aware of the treatment they are in (in the subjects' case, this is part of complete disclosure of the nature of the decision process, and is essential for the tests described herein). However, the subjects do not interact with the experimenter while they make decisions at their computer terminals, and the subjects do not know a priori the theoretical expectations

[7]In the actual experiment we described the decision situation to subjects in neutral, abstract terms. In particular, we referred to active types as *active* and to those who decided to vote as subjects who choose to be *available.* This removes a potential contaminating effect of "tipping off" the subjects about the kind of behavior that is somehow expected or appropriate.

or hypotheses about behavior in each treatment. Therefore, the danger of experimenter effects is minimal and the danger that subjects skew the results in favor of the theoretical expectations is also minimal.

We conducted two sessions of the experiment in computer labs at Northwestern University. Subjects were Northwestern undergraduates recruited from the Management and Organizations subject pool, undergraduate social science classes, and computer labs. Subjects were not selected to have any specialized training in game theory, political science, or economics. Sessions lasted for about 90 minutes and consisted of 90-100 rounds. For each subject, five rounds were selected at random at the end of the experiment and the subject was paid the sum total of her earnings in dollars from those rounds, times .04 (session 1) or .07 (session 2). Participants earned about \$25 on average for their session, with a minimum payment of \$15 up to a maximum of about \$50. Subjects were paid privately in cash at the end of the session so that a subject and the experimenter knew that subject's payment.

Session 1 had 18 participants and session 2 had 11. Each session began with an instruction period to familiarize the participants with the decision problem, computer software, random matching and sequence of decisions. The computer software displayed the payoff table (Table 1) with the experimental parameters, information about the subject's role and the number of subjects in each role in the group in a given round, and the entire history of the subject's own results. All decisions were made in private at computer terminals and all interaction among subjects took place anonymously at computers.

Session 1 consisted of 6 treatments of 15 rounds each for a total of 90 rounds. Session 2 consisted of 8 treatments of 10 or 15 rounds each for a total of 100 rounds.[8] The following table lists the values of n_β that were used with each combination of n_A and n_B in the experiment (number of rounds in which that value was used in parentheses). Recall that $n = n_A + n_B$ is the number of participants in each group.

[8] A design glitch in session 2 occurred after 18 rounds of the experiment. 18 rounds took place before the glitch, 15 in one treatment and 3 in another. Only 10 of 11 subjects were used in these rounds. In total, therefore, session 2 had 85 rounds with 11 subjects and 18 rounds with 10 subjects for a total of 103 rounds.

		n_A					
		2	3	5	6	8	9
	1	–	–	1(25)	–	–	–
	2	–	1(3)	–	–	–	–
	3	2(15)	–	–	–	1(15) 3(10)	–
n_B	4	–	–	1(40) 4(10)	–	–	–
	5	–	–	–	1(15) 5(15)	–	–
	8	–	–	–	–	–	1(15) 3(15) 7(15)

Table 2. Experiment design. Entries list number of active B types in group, for each possible combination of A and B types (no. of rounds for which the configuration was used in parentheses).

Therefore, the possible values of n_β were 1, 2, 3, 4, 5, and 7. Note that in all but 15 rounds, groups had more A voters than B voters ($n_A > n_B$). This ensures that option A maximizes the sum of payoffs received by the n members of a group, even though B maximizes the payoff of eligible voters and of B types collectively. Note also that given the cost of participation in our design, the cost of voting ($c = .1$) outweighs the maximum expected monetary benefit ($\frac{x}{2} = .125$) from voting, unless $n_\beta = 1$.

5 Results

5.1 Individual Behavior

Our most important results deal with the relationship between the probability a voter votes for each alternative as a function of the probability a vote is pivotal. In the absence of ethical expressive voters, i.e., all voters are a combination of selfish types and ethical instrumental types, we showed above that the incentive to vote for either alternative is decreasing in the probability a vote is pivotal. In contrast if all voters are ethical expressive types the incentive to vote for A is either constant or increasing in the probability a vote is pivotal while the probability of voting for B is strictly decreasing.

The following table presents multinomial logit[9] results from the data aggregated from all sessions. The standard errors are clustered by subject to reflect the fact that observations from a particular individual cannot be assumed to be independent. The baseline category is not voting. Therefore, coefficients for option A (the relatively even split of group gains) reflect the effect of each variable on the probability of voting for A as opposed to not voting, and coefficients for option B (the relatively lopsided split of gains, beneficial to B types) reflect the effect the effect of each variable on B as opposed to not voting.

Covariate	Parameter estimate	Clustered SE
Pr(Vote for A)		
A types	.289***	.116
B types	-.653***	.234
Pivot probability	-.05	.948
Round	-.014	.010
constant	1.134	1.705
Pr(Vote for B)		
A types	.090	.093
B types	-.139	.147
Pivot probability	1.899***	.539
Round	-.009	.008
constant	-.963	.848

Table 3. Effect of group characteristics on individual vote choice.
571 observations; Standard errors adjusted for 29 clusters
Note: * indicates $p < .10$, ** indicates $p < .05$, *** indicates $p < .01$

The results display a significant difference in the probability of voting for alternative A and B as a function of pivot probabilities. The probability a voter votes for B is significantly affected by pivot probabilities in the direction predicted by all three models. As the pivot probability increases the probability of voting for B increases. This can be seen by the estimated coefficient on Pivot probability in table 1. This coefficient is estimated at 1.899 (the marginal effect of removing one Active B type on the probability

[9] Hausman and Small-Hsiao tests of the Independence of Irrelevant Alternatives assumption cannot reject the null hypothesis that IIA is satisfied. Essentially, this reflects that no two choices are perceived as close substitutes for each other. In any case, multinomial probit results (which do not depend on IIA) reflect very similar effects.

of voting for B is .390) and is statistically significant at the level of 1%. In contrast changes in pivot probability have a statistically insignificant effect on the probability a voter votes for A. Such a result provides compelling support for the claim that the behavior of some voters is consistent with ethical expressive model.

Furthermore, the decision to vote for the selfish alternative B is unaffected by the welfare effects of this choice on the rest of the group (more precisely: these welfare effects are indistinguishable from zero statistically), while the decision to vote for the ethical alternative A is affected by the welfare effects.[10] These findings are reflected in the parameter estimates for the number of A types and number of B types in the group.

The results do not show an experience effect on either ethical or selfish voting. This is reflected in the Round variable, which indexes the round of the session in which a decision occurred. The effect is insignificant for both options relative to abstention.

5.2 Individual Behavior and Collective Choices

In this subsection we show that the individual-level effects predicted by the theory and identified above matter for the choices that groups make in the experiment, i.e. that the individual effects are economically relevant. The results show that pivot probability causes changes in group choices as well as individual ones, so that the preference effect is relevant at the social as well as individual level.

The table below presents logit results on the effect of group characteristics on the probability that the group choice was option B rather than option A. Because the variance of the observed outcome could change with the group characteristics, we report heteroskedasticity-robust standard errors.

[10] In principle, "altruistic" motivations in which individuals place some small positive weight on the welfare of others may make both participation and ethical voting rational from an instrumental point of view (Jankowski (2002), Edlin, Gelman, and Kaplan (2006)). Even if pivot probability is very small, for altruists the total utility difference between two alternatives can be large in a large electorate because of effect of the result on many other people. Linear utilities with constant altruism weights map straightforwardly into a structural model that is estimable in our data. This is a special case of the ethical instrumental model presented above. Multinomial logit estimates show that for this model to explain our experimental data, Active B types must value an A type's payoff as much as they value their own, and must value an inactive B voter's payoff *four times* as much as they value their own. These implausible findings show that a simple model of entirely ethical but consequentialist voters will not fare well in our data.

Covariate	Parameter estimate	Het. robust SE
Pr(Choice is B)		
A types	-.175	.109
B types	.377***	.132
Pivot Probability	.805*	.471
Random choice	-.666*	.270
Round	.008	.007
constant	-1.030	.713

Table 5. Effect of group characteristics on group choice.
281 observations; Heteroskedasticity-robust standard errors
Note: * indicates $p < .10$, ** indicates $p < .05$, *** indicates $p < .01$

The preference effect of pivot probability on the selfish choice is positive and significant — as pivot probability increases, the probability of a selfish choice for the increases. The marginal effect of pivot probability on the probability that the group choice is B is .200. The individual-level moral bias induced by pivot probability affects group choices as well.

Beyond this, the group choice is sensitive to the collective benefit for B types. As the number of B types increases, and their collective welfare gain of option B over option A does as well, the probability of B as the group choice increases rapidly.[11] On the other hand, while the estimate shows that an increase in the number of A types raises the likelihood of A, this effect is not significant.

The negative effect of Random Choice on the probability of B reflects that the selfish choice was more likely to be implemented by a subject in the experiment (roughly 59% of these cases) than by a random draw when the selected voter was unavailable (roughly 47% of these cases).

As with individual-level choices, group-level choices are not significantly affected by the experience of participants. This is reflected in the insignificant effect of Round in the session on the probability the group choice is B.

[11] At the group level, the number of B types and the number of Active B types have correlation of about 0.51. Therefore, with about 280 observations it is reasonable to assume that the parameter estimates on these variables are not excessively unreliable because of multi-collinearity.

6 Conclusion

In this paper we have provide experimental confirmation in support of an ethical expressive model of voting. In our experiment groups must choose between two options — an "ethical" option with a relatively broad distribution of monetary payments, and a "selfish" option with a more narrow distribution but favoring the voters themselves. Our design allows us to manipulate the collective welfare of each option and, most importantly, the pivot probability of individual voters. Therefore, we can control this crucial variable rather than leaving it as an endogenous choice.

The results from the experiment strongly suggest a preference effect of electorate size. As pivot probability declined, the broad distribution of gains in the group was significantly more likely to be chosen at both the individual and group levels. This points to a self-selection of other-regarding agents into the voting population as the electorate grows.

Our results also provide evidence that instrumentally rational voters (selfish or altruistic) are likely to abstain even for very small costs of voting in a large election. Voters who do choose to participate as pivot probability declines will be disproportionately expressive rather than instrumental in their motivation for voting. In our experiment this led to a disproportionate fraction of voters supporting an alternative consistent with an "other-regarding" ethical disposition. If this effect holds then large elections with voluntary participation would produce a sort of "benevolence amplification" in translating preferences in the electorate into election outcomes.

We should emphasize that while we call this effect benevolence amplification we are not suggesting that this is necessarily a normatively good property. Indeed, as we stated in the introduction, an argument for elections is that they choose policies that are broadly representative of the potential electorate. Our results indicate that this need not be the case.

7 Appendix

In this appendix we show the behavior of ethical expressive voters under all possible parameter values for c and d.With some algebra we can summarize the results into six cases

$d > 2c$	$n_\beta < \frac{x}{d}$	$n_\beta > \frac{x}{d}$
$x > d$	Vote B	Vote A
$x < d$	-	Vote A

$2c > d > c$	$\frac{x}{2c} > n_\beta$	$\frac{x}{2(d-c)} > n_\beta > \frac{x}{2c}$	$n_\beta > \frac{x}{2(d-c)}$
$x > 2c$	Vote B	abstain	Vote A
$x < 2c$	-	abstain	Vote A

$c > d > 0$	$\frac{x}{2c} > n_\beta$	$n_\beta > \frac{x}{2c}$
$x > 2c$	Vote B	abstain
$x < 2c$	-	abstain

Case 1 vote B vote A (d>2c and d<x)
Case 1* vote A (d>2c and d>=x)
Case 2 vote B abstain Vote A (2c>d>c and 2c<x)
Case 2* abstain Vote A (2c>d>c and 2c>x)
Case 3 Vote B abstain (d<c and x>2c)
Case 3* abstain (d<c and x<2c)

7.1 B as the Ethical Alternative and ethical instrumental expressive

Ethical expressive types get the same payoffs as selfish voters plus a payoff of $d > 0$ by voting for option A. The payoff to this type for voting for option A is

$$\frac{1}{n_\beta}(1+\delta) + \left(1 - \frac{1}{n_\beta}\right)(1 + q^* x + (1-q^*)\delta) + d - c$$

while the payoff for voting for option B is

$$\frac{1}{n_\beta}(1+x) + \left(1 - \frac{1}{n_\beta}\right)(1 + q^* x + (1-q^*)\delta) - c$$

Conditional on voting, ethical expressive voters prefer to vote for A over B if

$$d \geq \frac{x-\delta}{n_\beta}.$$

So, conditional on voting, as the probability of being pivotal decreases the incentive for an ethical expressive type to vote for A increases.[12] Note that this is in contrast to both the selfish and ethical instrumental models where pivot probabilities does not impact the choice between A and B.

[12] As the pivot probability increases the set of pairs (d, δ) that satisfy the equation above decreases (by inclusion).

Voters with $d \geq \frac{x-\delta}{n_\beta}$ prefer to vote for A rather than abstain if and only if

$$\begin{aligned} & \frac{1}{n_\beta}(1+\delta) + \left(1 - \frac{1}{n_\beta}\right)(1 + q^* x + (1-q^*)\delta) + d - c \\ \geq \; & \frac{1}{n_\beta}(1 + \frac{x+\delta}{2}) + \left(1 - \frac{1}{n_\beta}\right)(1 + q^* x + (1-q^*)\delta) \end{aligned}$$

or

$$d - c \geq \frac{x-\delta}{2n_\beta}.$$

Voters with $d < \frac{x-\delta}{n_\beta}$ prefer to vote for B rather than abstain if and only if

$$\begin{aligned} & \frac{1}{n_\beta}(1+x) + \left(1 - \frac{1}{n_\beta}\right)(1 + q^* x + (1-q^*)\delta) - c \\ \geq \; & \frac{1}{n_\beta}(1 + \frac{x+\delta}{2}) + \left(1 - \frac{1}{n_\beta}\right)(1 + q^* x + (1-q^*)\delta) \end{aligned}$$

or

$$\frac{x-\delta}{2n_\beta} \geq c.$$

Previous experimental work (Palfrey and Levine) claim to find support for the selfish voter model. Their results find that turnout in laboratory experiments conforms with the comparative statics predicted by the selfish model e.g., turnout is decreasing as the size of the electorate increases (and therefore pivot probabilities decrease). In this section we show that it is difficult to differentiate between the selfish and ethical expressive model when the alternative that is favored by selfish voters is also perceived to be the ethical alternative. We show that the only difference between the two models is in the level of turnout predicted. In the ethical expressive model turnout does not go to zero as pivot probabilities get small.

Suppose that ethical expressive types get a payoff of $\delta > 0$ when alternative B is chosen and a payoff of $d > 0$ by voting for option B. It is obvious that such voters will never vote for option A. The payoff for voting for option B is

$$\frac{1}{n_\beta}(1 + x + \delta) + \left(1 - \frac{1}{n_\beta}\right)(1 + q^*(x+\delta)) + d - c$$

Subjects prefer to vote for B rather than abstain if and only if

$$\begin{aligned} & \frac{1}{n_\beta}(1+x+\delta)+\left(1-\frac{1}{n_\beta}\right)(1+q^*(x+\delta))+d-c \\ \geq \; & \frac{1}{n_\beta}(1+\frac{x+\delta}{2})+\left(1-\frac{1}{n_\beta}\right)(1+q^*(x+\delta)) \end{aligned}$$

or

$$d-c \geq -\frac{x+\delta}{2n_\beta}.$$

So, if $d-c>0$ the subject votes while if $d-c<0$ then the probability of voting is decreasing in n_β. It follows that turnout is decreasing as pivot probabilities decrease but reaches a lower bound. Note that Palfrey and Levine observe in their experiments that turnout levels in elections with low pivot probabilities seem to be bounded significantly above zero.