Moral Bias in Large Elections: Theory and Experimental Evidence*

Timothy Feddersen[†] Sean Gailmard[‡] Alvaro Sandroni[§]

July 2007

*Thanks to Becky Morton and Jasjeet Sekhon for many helpful conversations; seminar participants at Dartmouth, Essex, Florida State, Oxford, and Penn; panelists at the 2006 annual meeting of the Midwest Political Science Association and the 2007 Society for the Promotion of Economic Theory conference; and Jenny Cheng for research assistance.

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

[‡]Travers Department of Political Science, UC Berkeley, 210 Barrows Hall, Berkeley, California 94720.

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

Abstract

We provide support for the claim that large elections may exhibit a *moral bias*, i.e., controlling for the distribution of preferences within the electorate, alternatives understood by voters to be morally superior are more likely to win in large elections than in small ones. 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.

1 Introduction

A central normative 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. 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 each person takes an *instrumental* approach to elections and votes for the option they prefer. In this simple example, everyone in the population would vote for B and it would win the election. Hence, standard theory predicts that elections (without asymmetric information and costs to vote) will produce outcomes that are representative of the preferences of the population.[1] Instrumental models of voting include voters whose preferences over outcomes are defined by the material payoffs that would result from each outcome but also include altruistic, inequality-averse, inefficiency-averse and other types as well.[2]

The standard model has trouble explaining turnout in large elections with costs to vote.[3] Riker and Ordeshook (1968), Tullock (1971), Brennan and Lomansky (1993), and Feddersen and Sandroni (2006a) among others have proposed *expressive* theories of voting to explain turnout in large elections. In expressive voting models individuals are motivated to vote, not out of a desire to directly impact the election outcome, but out of a sense of civic obligation or a desire to act ethically by supporting morally appealing causes or candidates. Brennan and Lomansky (1993) among others show that expressive voting models may give radically different predictions about the relationship between citizen preferences and election outcomes.

To see the difference between instrumental and expressive voting models assume, in the

[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, Levine and Palfrey 2007) and may bias election results in favor of those with lower costs to vote.

[2] See Jankowski (2002) and Edlin, Gelman and Kaplan (2006) for a discussion of altruism. See Fehr and Schmidt (1999) and Bolton and Okenfels (2000) on inequality aversion. See Charness and Rabin (2002) on inefficiency aversion.

[3] See Levine and Palfrey (2007) for an argument that standard models can explain turnout in large elections.

example above, that B gives higher material benefit than A to each voter in an electorate, but all agree 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 instrumental voting models agents will condition their vote choice on the event their vote is pivotal and all vote for option B.

Now consider agents who get a positive subjective payoff for taking actions they believe to be ethical.[4] This payoff is received as a function of what action the agent chooses and does not depend upon the outcome of the election. This subjective payoff is in addition to the usual instrumental payoffs.[5] We call *ethical expressive* voters those who receive a payoff $d > 0$ from taking an action they believe to be ethical.[6] Suppose the electorate consists entirely of ethical expressive voters and that the payoff for acting ethically is small enough so that all voters nevertheless prefer B to A. That is, if each voter were made a dictator they would vote for B. In large elections the probability a vote is pivotal is typically very small. When pivot probabilities are small, the voter's choice between A and B is very unlikely to have an impact on the outcome of the election. We say that such choices are *essentially hypothetical*.[7] 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, all vote for A rather than B. So, even a very small payoff to act ethically may have a large impact on election outcomes.

When there is no cost to vote, the impact of ethical expressive types is small if a large

[4]This subjective payoff can be thought of as a "warm-glow" payoff. See Andreoni (2006) for a review of the literature on warm-glow giving.

[5]See Harsanyi (1977), Coate and Conlin (2004), Feddersen (2004) and Feddersen and Sandroni (2006a, b) for examples of ethical voter models.

[6]This does not mean that such voters vote for an ethical alternative only that they obtain an additional positive payoff if they do irrespective of the election outcome.

[7]A choice that has no impact on the actual outcome of the election may reasonably be called a *hypothetical choice*. See Cummins et al (1997) and Holt and Laury (1997) for a discussion of hypothetical choices and differences in behavior compared to actual choices.

fraction of the electorate consists of instrumental voters. In that case all the instrumental types would vote for B and B would still be the outcome. However, when small costs to vote are introduced, the impact of ethical expressive types is amplified. This follows because instrumental agents become more likely to abstain as pivot probabilities decline. But ethical expressive types vote for A when pivot probabilities are small. Thus, in large elections we may expect the electorate to consist disproportionately of ethical expressive types.

The example above represents an extreme illustration of what we call *moral bias* in large elections i.e., controlling for the distribution of preferences within the electorate, alternatives understood by voters to be morally superior are more likely to win in large elections than in small ones. Moral bias occurs because of two effects. First, as pivot probabilities decrease, vote choice becomes more hypothetical and ethical expressive voters become more likely to vote for the alternative they understand to be ethically superior. We call this a *preference effect* of pivot probabilities. Second, in the presence of costs to vote, a decrease in pivot probabilities also reduces the incentive for instrumental types to vote and may increase the incentive of ethical expressive types to vote. As a consequence the electorate consists disproportionately of ethical expressive voters. We call this a *turnout effect*.

Tullock (1971), Brennan and Lomansky (1993) and others have observed that expressive voting can theoretically produce outcomes unrepresentative of preferences in the electorate. However, it remains an open question whether large elections really create the incentives leading to moral bias. Survey and election data have trouble resolving voter motivations (Sears and Lau 1983). Results from experiments to date are also ambiguous. Fischer (1995) and Carter and Guerette et al. (1990) find weak support, while Tyran (2004) finds no support for expressive voting. Levine and Palfrey (2007) provide experimental results that support the claim that standard instrumental models of voting with error-prone behavior can explain voting behavior in large elections.

The contribution of this paper is to provide experimental evidence for both the preference and turnout effects consistent with expressive but not instrumental voting models. We construct an experiment in which a group chooses between one of two alternatives: A or B. The population is subdivided into A types who get a high payment if A is chosen and nothing if B is the outcome. B types get a high payment if B is chosen and a smaller payment if A is selected. A majority of the population are A types. Alternative A maximizes the sum of payments, gives nearly equal payments to everyone and maximizes the minimum payment. For these reasons we call A the *ethical* alternative.

Experimental elections complicate analyzing the effects of changes in pivot probabilities because of a multiplicity of equilibria.[8] Instead of this approach we simulate an election with costly voting with a decision mechanism in which pivot probability is controlled directly as a treatment variable. In our experiment design, a subset of individuals are designated as *active* and may either vote for A or B at a 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 A and B are chosen with equal probability. 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 $1/n$.[9] In our experiment only B types are active, and by changing the number of active B types we can change pivot probabilities. This allows us to focus on the impact of changes in pivot probabilities on the incentives for people to vote against their material interests. By varying the ratio of A to B types we can manipulate the degree to which A might be perceived as a morally superior alternative to B according to ethical theories such as utilitarianism.

In section 2 of the paper we formally define the payoffs for instrumental and ethical expressive types and derive optimal behavior for each as a function of pivot probabilities. Under the assumption that ethical expressive motivations substantially affect the behavior of at least some individuals, our theoretical results imply that as pivot probabilities decline: (1) the

[8] See Levine and Palfrey (2007) and Duffy and Tavits (2006).
[9] See Plott and Grether (1979) for an example of such an approach.

probability a subject votes for B rather than abstains should decline; (2) the probability a subject votes for A rather than abstains should not decline; (3) the probability a subject votes for A rather than B should be increasing; and (4) the probability A is the collective outcome is increasing. Predictions 1 and 2 are consistent with the hypothesized turnout effect. Prediction 3 is consistent with a preference effect. Prediction 4 is consistent with moral bias. Our experimental results are consistent with all these hypotheses. Finally, a quantal response model (see Levine and Palfrey 2007) does not account for the data as well as the expressive voting model.

The paper is organized as follows. In section 2 we define a formal model of the experimental design and derive behavioral predictions. In section 3 we describe how the experimental design was implemented. In section 4 we describe the results of the experiments and in section 5 we provide a brief conclusion.

2 Behavioral Predictions

Before we describe the experimental design in detail it will be helpful to describe the formal 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 $n_{\text{-}\beta}$ be the number of inactive B types so that $n_\beta + n_{\text{-}\beta} = n_B$. 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.

Monetary rewards are given below:

Table 1. 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. 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 $\frac{1}{2} > x > 2c > 0$ and $n_A > n_B$. These assumptions ensure that alternative A minimizes inequality in terms of monetary rewards, maximizes the sum of monetary rewards, maximizes the minimum reward, and gives a higher monetary reward to a majority of the group. For these reasons we say that A is the *ethical outcome*.[10]

2.1 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 the incentives to vote in this model 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

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

[10]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 Appendix A.

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.

2.2 Ethical Instrumental Types

Voters may depart from selfish behavior if they take into account the monetary rewards of others. A large literature in experimental economics and game theory suggests that such considerations are important (e.g. Fehr and Schmidt 1999; Bolton and Ockefels 2000; Charness and Rabin 2002; see Camerer 2003 a partial review). We define *ethical instrumental* voters as those who prefer option A to option B. Ethical instrumental agents, like selfish (instrumental) types care about their vote only insofar as it affects the outcome of the election.

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.

2.3 Ethical Expressive Types

We call agents who get a payoff simply by voting for option A whether or not their vote is pivotal *ethical expressive types.* The key concept is that ethical expressive voters get a payoff for taking an action they determine to be ethical, independent of the consequences of that action. For other examples and an application to a new theory of electoral participation, see Feddersen and Sandroni (2006a, 2006b).[11]

[11] Riker and Ordeshook (1968), Brennan and Buchanan (1984), Uhlaner (1989), and Scheussler (2000) all present non-instrumental theories of voter participation, and Brennan and Buchanan and Scheussler focus specifically on types of expressive benefits. However, none of these theories explores the interaction of instrumental and expressive motivations as we do, or explore the effects of pivot probability on participation, vote choice, or collective decisions. Another perspective related to ours is Andreoni (1990), who models a "warm glow" from the act of giving *per se* in the context of public goods.

Ethical expressive types get the same payoffs as selfish voters plus a payoff of $d > c$ by voting for option A.[12] The payoff to this type for voting for option A is

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

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. Note that this is in contrast to both the selfish and ethical instrumental models where pivot probabilities do 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

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

[12] The assumptions that $d > c$ and that A is the ethical option (for which the expressive payoff d occurs) ensure that the behavior of the ethical expressive voters is different qualitatively from the behavior of the selfish voters. We relax these assumptions in Appendices A and B.

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 (see Appendix A) that the behavior of ethical expressive types is as follows. Consider three different cases: $d \geq x$ (d large); $x > d > 2c$ (d intermediate); and $2c > d$ (d low). 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. 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 when d is small ($\frac{2(d-c)}{x} > \frac{1}{n_\beta}$).

The key point is that ethical expressive types behave much differently from either selfish or instrumental types. Therefore, their presence has an important effect on the relationship between pivot probability and the collective choice in the election. Specifically, ethical expressive 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 is the *preference effect* of pivot probability alluded to in the introduction. 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* because it does not 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, which is not affected by pivot probability for ethical expressive agents, 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 non-decreasing or even increasing as pivot probabilities decrease whereas in the latter cases the incentive to vote is decreasing as pivot probabilities decrease. This is the *turnout effect* of pivot probability alluded to in the introduction.

The three different models can be unified into a single framework. For purposes of exposition we specified deterministic models above. Given that these are decision-theoretic models it is natural to include a stochastic element to account for apparently random behavior typically observed in experiments.[13] The payoffs for each choice in a unified model are specified as follows. The payoff of voting for alternative A is

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

The payoff of voting for alternative B is

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

The payoff of not voting is

$$\pi_\phi(n_\beta, x) = \frac{1}{n_\beta}\left(1 + \frac{x+\delta}{2}\right) + \left(1 - \frac{1}{n_\beta}\right)(1 + \delta + q^*(x-\delta)) - \varepsilon.$$

Note that the differences between these payoffs are all linear in δ and d. Furthermore, as we will explain below, n_β, x and c are all control variables in the experimental design. As above q^*(the probability alternative B is selected when the decision-maker's vote is not pivotal) drops out of the differences. Given this and the standard assumption that ε follows a type I extreme value distribution, the choice among the three alternatives can be modeled in a standard multinomial logit model.

[13]See McKelvey and Palfrey (1995).

2.4 Hypotheses

From this analysis one can predict the effects of pivot probability on participation and vote choice by agents, and therefore the effects of pivot probability on the probability that each option is chosen for the group.

Specifically, suppose all agents have payoffs as specified in the unified model above. Then, if $d > c$ the model predicts that as pivot probability declines,

1. Agents are less likely to vote for B rather than abstain. Any instrumental types who prefer option B are less likely to vote as pivot probability declines. This hypothesis is also implied by the ethical expressive model if the payoff from acting ethically is small enough.[14]

2. The probability an agent votes for A rather than abstains is nondecreasing. This is because of the presence of ethical expressive voters who support option A. Ethical expressives' incentive to vote for A does not result from the possible effect of a vote on the election outcome, so is not sensitive to decreasing probability of such an effect. On the other hand their incentive to vote against A is decreasing in pivot probability. Thus ethical expressives have either constant or increasing probability of voting for A as pivot probability decreases.

3. The probability a subject votes for A rather than B is increasing. Conditional on voting, agents are more likely to select the ethical option A. This follows from the preference effect. When pivot probability declines vote choice is closer to a hypothetical choice, so the non-instrumental component of utility weighs more heavily.

4. The group is more likely to select the ethical option A. This follows from both the preference and turnout effects.

[14] If the payoff from acting ethically is large then ethical expressives already have 0 probability of voting for B. See Appendix A.

3 Experiment Design

The experiment design described below allows us to test the predictions from Section 2 in a controlled setting. This control is crucial for our study for two reasons. First it allows us to directly manipulate pivot probability, the key causal variable in our theory. As we elaborate below, we take extra steps to control pivot probability itself, rather than simply control group size and rely on equilibrium reasoning to translate this into pivot probability. Second, experimental control allows us to induce specific monetary values for the options facing the group. Even though we do not fully control preferences, control over monetary payoffs still allows us to effectively determine which option is selfishly beneficial to voters and which option is ethical in several respects.

The experiment was conducted in a sequence of rounds. A round in turn consists 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, n_β, and as we elaborate below, the payoffs for options A and B to every group member. 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, $c = 0.10$ denotes the participation cost and $x = 0.25$ 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.[15]

We conducted four 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. Sixty-one subjects participated across the four sessions, with subjects per session ranging from 9 to 24. 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 not visible to other subjects and all interaction among subjects

[15]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.

took place anonymously at computers.

Sessions lasted for about 90-100 minutes, consisted of 90-150 rounds, and contained six to eight distinct treatments.[16] 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 0.04. 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.

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. The treatments were chosen primarily to maximize the range of possible values of n_β (and therefore pivot probability) given the number of subjects in each session and $n_A > n_B$ while still varying the ratio of n_A to n_B. A consequence of this is that the design is incomplete in a factorial sense. Note that for most (n_A, n_B) pairs, n_β ranges roughly as much as possible with high contrast between treatments. Treatments were ordered in a crossover fashion within sessions and counterbalanced fashion across sessions (again, given the constraint of the number of subjects) to control for order and testing/learning effects.

[16] A software 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_B \ n_A	2	3	5	6	7	8	9	13
1	–	1(10)	1(25)	–	–	–	–	–
2	2(10)	1(10)	–	–	–	–	–	–
3	2(15)	–	1(10) 3(10)	2(15) 3(15)	– –	1(15) 3(10)	–	–
4	1(10) 4(10)	–	1(40) 2(15) 3(15) 4(35)	–	–	–	–	–
5	–	–	–	1(15) 5(15)	1(10) 4(10) 5(10)	–	–	–
8	–	–	–	–	–	–	1(15) 3(15) 7(15)	–
11	–	–	–	–	–	–	–	2(15) 11(20)

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 (160 rounds), 2 (70 rounds), 3 (65 rounds), 4 (55 rounds), 5 (25 rounds), 7 (15 rounds), and 11 (20 rounds). Note that in almost all 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 ($c = 0.10$) outweighs the maximum expected monetary benefit ($\frac{x}{2} = 0.125$) from voting, unless $n_\beta = 1$.[17]

[17] 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 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.

4 Results

4.1 Individual Behavior

The key effects of pivot probability operate at the individual level. Thus we begin the analysis there. As noted above, for selfish types and ethical instrumental types, the incentive to vote for either alternative decreases as pivot probability decreases. For ethical expressive types, the incentive to vote for A is either constant or increasing in pivot probability while the probability of voting for B is weakly decreasing (strictly so if the payoff from ethical voting is small enough to give these types an incentive to vote B when pivot probability is 1).

Figure 1 presents graphical evidence on these points. It displays mean individual vote choice decisions as a function of pivot probability, across treatments with more than one pivot probability value (thus, group composition is held constant in each panel). Hypothesis 1 implies that the dashed "selfish voting" line should slope up; hypothesis 2 implies that the solid "ethical voting" line should be flat or downward sloping.[18] Almost all panels are consistent with this. The main exception is the $(N_A, N_B) = (6, 3)$ panel, which is also the weakest (lowest contrast) treatment in the figure in that the number of active B votes varies only from 2 to 3. This panel supports hypothesis 1 but not hypothesis 2. In addition, in the $(N_A, N_B) = (7, 5)$ panel mean vote choice is nearly constant with pivot probability, which is not consistent with hypothesis 1. These exceptions aside, the raw data reflected in these figures is generally consistent with the predictions.

[Figure 1 here]

Another simple and compact way to assess the effect of pivot probability on participation and vote choice is through statistical models. The following table presents multinomial logit[19] results from the data aggregated from all sessions. The standard errors are clustered by

[18] The slope of the dotted Abstention line is not restricted by the theory.

[19] 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.

subject to reflect the fact that observations from a particular individual cannot be assumed to be independent. The baseline category is abstention. 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.

The results reveal a significant difference in the effect of pivot probability on the probability of voting for each alternative rather than abstaining. Consistent with Hypothesis 1, the probability an agent votes for B rather than abstains is significantly affected by pivot probabilities in the direction predicted for all types of voters: 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 (the marginal effect of removing one Active B type on the probability of voting for B is .390) which is statistically significant at the level of 1%.[20]

Covariate	Parameter estimate	Clustered SE
Pr(Vote for A)		
Number of A types	0.210**	0.089
Number of B types	-0.349***	0.087
Pivot probability	0.524	0.377
Round	-0.001	0.003
constant	1.075*	0.604
Pr(Vote for B)		
Number of A types	0.082	0.052
Number of B types	-0.054	0.076
Pivot probability	1.372***	0.365
Round	-0.005	0.004
constant	-0.982**	0.458

Table 3. Effect of group characteristics on individual vote choice (multinomial logit; base category is Abstain).
1531 observations; Standard errors adjusted for 58 clusters
Note: * indicates $p < 0.10$, ** indicates $p < 0.05$, *** indicates $p < 0.01$

[20]It might be suspected that the number of B types is highly correlated with pivot probability, since the latter is isomorphic to the number of active B types. However, this correlation is less than 0.30. In any case, multicollinearity problems would typically be suspected of making intercorrelated variables appear falsely insignificant, not falsely significant.

In contrast, and consistent with Hypothesis 2, changes in pivot probability have a statistically insignificant effect on the probability a voter votes for A. The estimated effect is much smaller than on voting for B rather than abstaining. The propensity to vote for A does not decline when pivot probability declines. This supports the claim that the behavior of some voters is driven by expressive factors, and they tend to make choices on ethical grounds.

Beyond the implications for the hypotheses, the models also reveal interesting effects of group welfare on vote choice. The decision to vote for the selfish alternative B is unaffected by the utilitarian 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 utilitarian welfare effects. These findings are reflected in the parameter estimates for the number of A types and number of B types in the group. These variables determine the collective benefits of each option for the group. As the number of A types grows (holding fixed the number of B types), ethical considerations (particularly utilitarian ones) point more strongly in favor of alternative A.[21] As the number of B types grows, the pull of ethical considerations in favor of option A weakens. Selfish motivations in vote choice are necessarily insensitive to these considerations, and the results are consistent with the assumption of that some have selfish motivations.

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 and suggests that our findings above are not an artifact of inexperienced play.

These results are consistent with Hypotheses 1 and 2, but they may also appear consistent with purely instrumental, other-regarding motivations. In principle, "altruistic" motivations in which individuals place some small positive weight on the welfare of others may make

[21] Given fixed payoffs for options A and B, only some types of ethical considerations are affected by group composition in this way. For example, maximin considerations are unaffected by the size and composition of each group. Therefore, changing group composition need not be linked to all ethical considerations of expressive voters, and the ethical expressive model does not necessarily imply any particular effects of group composition on turnout and vote choice.

both participation and ethical voting rational from an instrumental point of view (Jankowski 2002; Edlin, Gelman, and Kaplan 2006; Fowler 2006). Even if the pivot probability is very small, for altruists the total utility difference between two alternatives can be large in a large electorate because of the effect of the result on many other people, thereby giving a greater incentive to vote. Given a decision among two alternatives, this is a special case of the ethical instrumental model presented above, and is not an adequate explanation for our data for two reasons. First, our design allows us to control the group size separately from pivot probability, and conditional on group size, decreasing pivot probability should still depress participation rates for these instrumental A voters. Yet the results show that abstention is not significantly more likely relative to voting for A as pivot probability declines, holding fixed the number of A and B types. Thus, while altruists may have higher turnout levels than selfish agents, their response to pivot probability should be the same. This is not consistent with our data. A possible counterargument is that these ethical instrumental voters' other-regarding benefits of voting are so high that even our lowest pivot probability does not depress their turnout. This leads to a second problem with this explanation for our findings. Linear utilities with constant altruism weights map straightforwardly into a structural model that is estimable in our data. 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 payoffs to other B types four times as much as they value their own. These implausible findings suggest that a model consisting only of instrumental voters does not fare well in our data.

In the ethical expressive model the expressive component of utility exerts a greater effect on decisions as pivot probability declines. This in turn implies that, conditional on turnout, voting for the ethical option grows more likely as pivot probability declines (Hypothesis 3). To test this hypothesis we model vote choice conditional on turnout. This is a binary choice for which we model the logit, with standard errors clustered by subject.

Covariate	Parameter estimate	Clustered SE
Pr(Vote for B)		
Number of A types	-0.144	0.090
Number of B types	0.312***	0.106
Pivot Probability	0.782*	0.465
Round	0.008*	0.004
constant	-0.944	0.729

Table 4. Vote choice conditional on turnout (logit).
625 observations; Standard errors adjusted for 58 clusters
Note: * indicates $p < 0.10$, ** indicates $p < 0.05$, *** indicates $p < 0.01$

Consistent with Hypothesis 3, a decline in pivot probability does make a vote for the ethical alternative more likely conditional on turnout. The result is significant at the 10% level, and the estimated marginal effect at the average is 0.19, a substantial response.

The number of each voter type captures the collective benefits of each option and measures one (utilitarian) aspect of ethical preferences. The number of B types in the group has a significant effect on the probability of voting for B. This effect is larger in absolute value (and in the opposite direction of) the effect of the number of A types on the probability of voting for B, which has a p-value of 0.11. The reason for this asymmetry is not obvious. One possibility is that agents find "excuses" to rationalize voting for the selfishly beneficial option, displaying a kind of self-serving perception of fairness found in bargaining games (e.g., Knez and Camerer 1995; Diermeier and Gailmard 2006).

In this case there is a statistically significant effect of the round of the session: agents are more likely to vote selfishly, conditional on turnout, in later rounds than earlier ones. This could be due to learning effects, or because of attempts to support cooperative play (widely dispersing benefits) that unravel as the session progresses. In any case, while these possibilities may make for interesting future explorations, the effect is not overwhelming in practical significance.

4.2 Individual Behavior and Collective Choices

The results above show the importance of ethical expressive considerations in making sense of individual behavior in our experiment. But this does not by itself demonstrate the importance of these considerations for understanding collective choices. The individual-level effects may be small or ethical expressive motivations so unusual as to be irrelevant for understanding the functioning of elections in collective decision-making. Since the ultimate importance of our analysis lies in how political institutions translate preferences into collective decisions in the presence of non-selfish, expressive agents, the question of group-level behavior is crucial and we turn to it now. The results show that pivot probability causes changes in group choices consistent with the presence of ethical expressive motivations, so that their effect is relevant at the social as well as individual level.

One simple way to demonstrate the effect of pivot probability and other group conditions on the collective choice is with statistical models of the group outcome. 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.

The most important feature of the results is that they support Hypothesis 4. The effect of pivot probability on the selfish choice is positive and significant. As pivot probability increases, the probability of a selfish choice for the group increases. The coefficient is significantly different from 0 at the 10% level. The marginal effect of pivot probability (at its average) on the probability that the group choice is B is .200, a sizeable effect. The individual-level effects carry over to group choices as well.

Covariate	Parameter estimate	Het. robust SE
Pr(Group choice is B)		
A types	-0.175	0.109
B types	0.377***	0.132
Pivot Probability	0.805*	0.471
Random choice	-0.666*	0.270
Round	0.008	0.007
constant	-1.030	0.713

Table 5. Effect of group characteristics on group choice (logit).
616 observations; Heteroskedasticity-robust standard errors
Note: * indicates $p < 0.10$, ** indicates $p < 0.05$, *** indicates $p < 0.01$

Beyond this, the group choice is sensitive to the collective benefit for B types, who benefit from the lopsided distribution of gains in the group. 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. 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. It is not apparent why this asymmetry exists.

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). We include this control variable in the model to ensure that the relationship between pivot probability and group choice is not driven entirely by the abstention of selfish voters as pivot probability declines (irrespective of the presence or behavior of ethical expressives), whose support for A is replaced by a coin flip.

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.

4.3 Quantal Response Model

Levine and Palfrey (2007) argue that quantal response equilibrium models with selfish instrumental voters can explain turnout and vote choice in large elections. Our experiment presents subjects with a decision situation, not a game, so in our unified model presented above, quantal response equilibrium collapses to quantal choice or random utility maximization (McFadden 1974), and a logit QRE is simply a logit random utility model.

Assume that d and δ are zero so that behavior conforms to a selfish quantal choice model. Three implications of this special case emerge:

1. The probability of abstention increases as pivot probabilities decline. This is because, assuming voting is a dominated strategy, the "error" of turning out becomes more costly as pivot probability declines.

2. Conditional on voting, the probability an agent votes for B is greater than the probability s/he votes for A. This is because, in quantal choice models, the probability of playing better strategies exceeds that of playing worse strategies.

3. Conditional on voting, the probability of voting for A increases and the probability of voting for B decreases as pivot probability declines. This is because the "error" of voting for A becomes less costly as pivot probability declines, as it is less likely to affect the outcome.

Only the first and second implications differentiate the selfish instrumental QRE model from the ethical expressive model, and as noted above they are not consistent with our experimental findings. Implication three follows from both the selfish instrumental and ethical expressive models, when a quantal choice component is included in each. It is supported in our data in the sense that agents are more likely to vote for A as pivot probabilities decline. Overall, then, our findings provide little support for a selfish instrumental QRE model to the extent that it differs from the ethical expressive model.

A major reason for the difference in our results and Levine and Palfrey's, which led them to conclude the selfish instrumental QRE is substantially supported, is that their experimental design suppresses ethical considerations. Groups essentially play a zero sum game in their design, so there is no difference between the ethical alternative and the selfishly beneficial one.

5 Conclusion

In this paper we provide experimental support for an ethical expressive model of voting. In our experiment groups must choose between two options — an “ethical” option with a relatively equal distribution of payoffs that maximizes total payoffs and the minimum payoff in the group, and a “selfish” option with a lopsided distribution favoring the voters themselves. Our design allows us to manipulate the distribution of payoffs from 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 variable.

The experimental results support the concept of moral bias in large elections: collective choices in elections systematically depart from individual preferences in the direction of moral considerations as pivot probabilities decline. The data is consistent with significant presence of ethical expressive types in the population. Moral bias then results for two distinct reasons: first, as the pivot probability declines, the choice of any agent that actually votes becomes closer to a hypothetical choice, in which case ethical considerations dominate selfish ones (the preference effect). Second, as pivot probability declines, instrumental voters are less likely to vote, while ethical expressive voters may continue to vote or switch from abstention to voting for the ethical option. Thus the ratio of ethical expressive voters to instrumental voters grows as pivot probability declines (the turnout effect).

Appendix A Optimal Behavior of Ethical Expressive Voters

In this appendix we present the optimal behavior of ethical expressive voters under all possible parameter values for c and d.

Case 1: $d > 2c$	$n_\beta < \frac{x}{d}$	$n_\beta > \frac{x}{d}$
$x > d$	Vote B	Vote A
$x < d$	-	Vote A

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

Case 3: $c > d > 0$	$\frac{x}{2c} > n_\beta$	$n_\beta > \frac{x}{2c}$
$x > 2c$	Vote B	abstain
$x < 2c$	-	abstain

The results can be grouped into six cases (one for each row of each table) of the relationship between pivot probability and vote choice. The cases list the sequence of optimal decisions for ethical expressive voters as pivot probability declines below the cutpoints in the tables. In all cases, the incentive to vote B is strongest for high pivot probabilities, the incentive to vote A is strongest for low pivot probabilities, and the incentive to abstain is highest for moderate pivot probabilities.

Case 1 ($d > 2c$ and $d < x$): vote B, vote A

Case 1* ($d > 2c$ and $d > x$): vote A

Case 2 ($2c > d > c$ and $2c < x$): vote B, abstain, vote A

Case 2* ($2c > d > c$ and $2c > x$): abstain, Vote A

Case 3 ($d < c$ and $x > 2c$): vote B, abstain

Case 3* ($d < c$ and $x < 2c$): abstain.

Appendix B Ethical Expressive Voting when Option B is the Ethical Alternative

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

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

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

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 (Levine and Palfrey 2006) shows some support for the selfish voter model in comparative statics. 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 Levine and Palfrey observe in their experiments that turnout levels in elections with low pivot probabilities seem to be bounded significantly above zero.

References

Andreoni, James (1990). Impure Altruism and Donations to Public Goods: A Theory of Warm-Glow. *Economic Journal* 100: 464-477.

Andreoni, James (2006). Philanthropy. In L-A Gerard-Veret, S-C Kolm, and J. Mercier, eds., *Handbook of Giving, Reciprocity, and Altruism.* New York: Elsevier.

Bolton, Gary and Axel Ockenfels (2000). ERC: A Theory of Equity, Reciprocity, and Competition. *American Economic Review* 90: 166-193.

Brennan, Geoffry and James Buchanan (1984). Voter Choice: Evaluating Political Alternatives. *American Behavioral Scientist* 28: 185–201.

Brennan, Geoffrey and Loren Lomasky (1993). *Democracy and Decision: The Pure Theory of Electoral Preference.* New York: Cambridge University Press.

Carter, John. and Stephen Guerette (1992). An Experimental Study of Expressive Voting. *Public Choice* 73: 251-260.

Cummins, R.G., S. Eliott, G.W. Harrison, and J. Murphy (1997). Are Hypothetical Referenda Incentive Compatible? *Journal of Political Economy* 103: 609-621.

Camerer, Colin (2003). *Behavioral Game Theory: Experiments in Strategic Interaction.* Princeton: Princeton University Press.

Charness, Gary, and Matthew Rabin (2002). Understanding Social Preferences with Simple Tests. *Quarterly Journal of Economics* 117: 817-69.

Coate, Stephen and Michael Conlin (2004). A Group Rule-Utilitarian Approach to Voter Turnout: Theory and Evidence. *American Economic Review* 94: 1476-1504.

Diermeier, Daniel and Sean Gailmard (2006). Self-interest, Inequality, and Entitlement in Majoritarian Decision-making. *Quarterly Journal of Political Science* 1(4).

Duffy, John and Margit Tavits (2006). Beliefs and Voting Decisions: A Test of the Pivotal Voter Model. Typescript, University of Pittsburgh.

Edlin, Aaron, Andrew Gelman, and Noah Kaplan (2005). Voting as a Rational Choice: The Effect of Preferences Regarding the Well-Being of Others. Typescript, University of California–Berkeley.

Feddersen, Timothy (2004). Rational Choice Theory and the Paradox of Not Voting. *Journal of Economic Perspectives* 18: 99-112.

Feddersen, Timothy and Alvaro Sandroni (2006a). A Theory of Participation in Elections. *American Economic Review.*

Feddersen, Timothy and Alvaro Sandroni (2006b). Ethical Voters and Costly Information Acquisition. *Quarterly Journal of Political Science* 1(3).

Fehr, Ernst and Klaus Schmidt (1999). A Theory of Fairness, Competition, and Co-operation. *Quarterly Journal of Economics* 114: 817-868.

Fischer, A.J. (1995). A Further Experimental Study of Expressive Voting. *Public Choice* 88: 171-180.

Fowler, James (2006). Altruism and Turnout. *Journal of Politics* 68.

Grether, David and Charles Plott (1979). Economic Theory and the Preference Reversal Phenomenon. *American Economic Review* 69: 623-638.

Harsanyi, John (1977). Rule Utilitarianism and Decision Theory. *Erkenntnis* 11: 25-53.

Holt, Charles and Susan Laury (1997). Theoretical Explanations of Treatment Effects in Voluntary Contribution Experiments. In Charles Plott and Vernon Smith, eds., *Handbook of Experimental Economics Results.* New York: Elsevier.

Jankowski, Richard (2002). Buying a Lottery Ticket to Help the Poor: Altruism, Civic Duty, and Self Interest in the Decision to Vote. *Rationality and Society* 14: 55-77.

Knez, Marc and Colin Camerer (1995). Social Comparison and Outside Options in Three-Player Ultimatum Game Experiments. *Games and Economic Behavior* 10: 65-94.

Levine, David and Thomas Palfrey (2007). The Paradox of Voter Participation? A Laboratory Study. *American Political Science Review.*

McFadden, Daniel (1974). Conditional Logit Analysis of Qualitative Choice Behavior. In P. Zarembka (ed.), *Frontiers of Econometrics.* Academic Press.

McKelvey, Richard and Thomas R. Palfrey (1995). Quantal Response Equilibrium in Normal Form Games. *Games and Economic Behavior* 10: 6-38.

Palfrey, Thomas and Howard Rosenthal (1983). A Strategic Calculus of Voting. *Public Choice* 43: 7–53.

Riker, William and Peter Ordeshook (1968). A Theory of the Calculus of Voting. *American Political Science Review* 62: 25-43.

Sears, David and Richard Lau (1983). Inducing Apparently Self-Interested Political Preferences. *American Journal of Political Science* 27: 223-252.

Scheussler, Alex (2000). *A Logic of Expressive Choice.* Princeton: Princeton University Press.

Tyran, Jean-Robert (2004). Voting when Money and Morals Conflict: An Experimental Test of Expressive Voting. *Journal of Public Economics* 88: 1645-1664.

Tullock, Gordon (1971). The Charity of the Uncharitable. *Western Economic Journal* 9: 379-392.

Uhlaner, Carol (1989). Relational Goods and Participation: Incorporating Sociability into a Theory of Rational Action. *Public Choice* 62: 255–285.

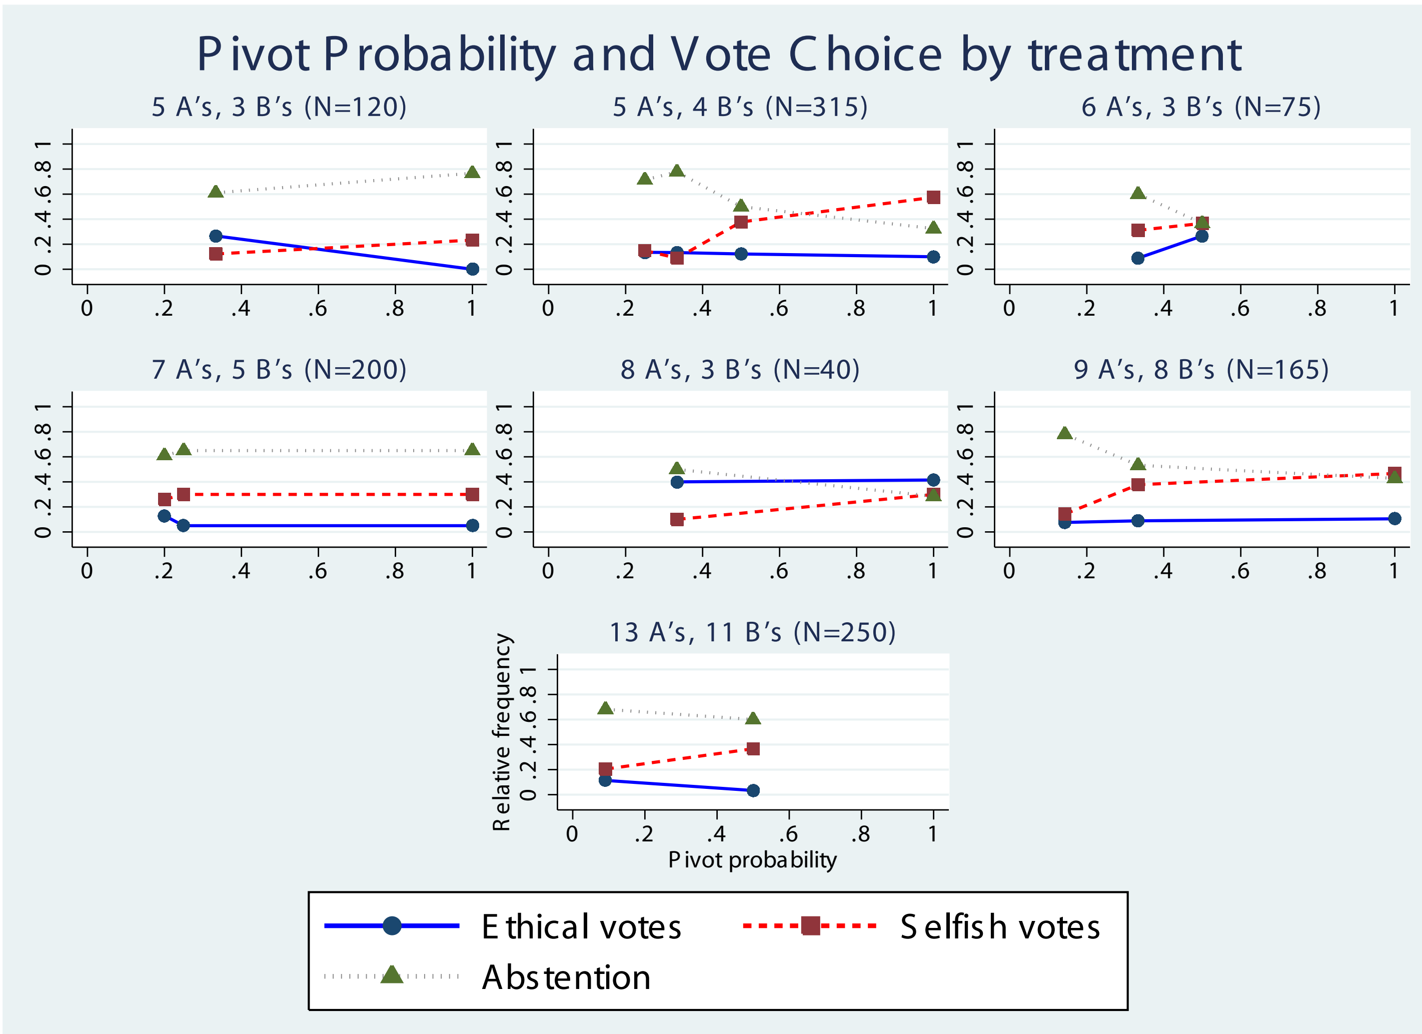

Figure 1:
Mean individual decisions from (N_A, N_B) pairs with more than one pivot probability value. N in each panel refers to the number of individual decisions from all active B subjects.