An Empirically-Grounded Theory of Voting Based on Prosocial Motives*

Dan Bernhardt† Laurent Bouton‡ Stefan Krasa§ Francesco Squintani¶

August 31, 2026

Abstract

We propose a model of costly voting in a canonical spatial setting that reconciles key empirical features of electoral turnout. We integrate citizens with prosocial voting motives who realize higher payoffs from voting when (i) they expect more fellow supporters to vote, (ii) such votes collectively matter more for winning probabilities, and (iii) they care more about which candidate wins. We establish existence of Nash equilibria satisfying an additional condition that ensures equilibria correspond to those in a finite setting with many voters. We derive how electoral primitives—e.g., voter ideology, voting costs, elite polarization, expected closeness—affect outcomes such as turnout, turnout rates, and winning margins. We relate the predictions to empirical findings both from existing studies and our original data work. With all three prosocial motive components, our model reconciles the empirical regularities qualitatively and quantitatively; while existing models fail on multiple dimensions.

*We gratefully acknowledge helpful comments and suggestions from Hans Gersbach, Georgy Egorov, seminar participants at Caltech, GMU, Kellogg, UCLA, Yale, Banco de Mexico, and conference participants at GLPET 2025, NBER 2025, and ETH Zürich. Maxime Cugnon de Sevricourt provided outstanding research assistance.

†Department of Economics, University of Illinois, danber@illinois.edu, Department of Economics, University of Warwick, m.d.bernhardt@warwick.ac.uk

‡Department of Economics, Georgetown University and NBER, lb910@georgetown.edu

§Department of Economics, University of Illinois, skrasa@illinois.edu

¶Department of Economics, University of Warwick, f.squintani@warwick.ac.uk

1 Introduction

Voters are at the heart of the democratic decision-making process. They influence policy directly through referenda and ballot initiatives, and indirectly by electing representatives and holding them accountable. In 2024 alone, about 70 nations—comprising roughly half of the world's population—held elections to choose their national leaders, including presidents and national assembly members. Beyond these national-level elections, citizens often participate in local elections to select state legislators, mayors, school board members, etc. Despite this, scholars remain divided on how to model voter behavior, and, as we show, all existing models deliver multiple predictions that are at odds with empirical evidence. This impairs our understanding of voting behavior and our ability to address important policy questions. Moreover, even when a political economy model does not explicitly focus on voters, accurately accounting for voter behavior is essential to capture voters' influence on other actors and dimensions of the democratic decision-making process.

The lack of consensus does not imply an absence of voting models. To the contrary, many model variants exist, including the so-called expressive voting, pivotal voting, abstention by indifference and/or alienation, and group-based voting models. However, as we show, these models fail to align with a number of empirical features (e.g., regarding voter turnout and win margins) that we document. For example, expressive voters are unrealistically insensitive to their electoral environment, while the tiny pivotal voting probabilities in large elections provide essentially no reason to vote, leading to the so-called "paradox of voting". Group-based models and models with abstention by indifference and/or alienation can deliver non-trivial turnout rates but predict various counter-factual voting patterns (e.g., a W-shaped pattern between ideology and abstention, as well as sharp declines in total turnout in lopsided elections absent counterfactually high levels of uncertainty about who will win).

We develop a tractable model of voting that we show can reconcile the core empirical features of voter turnout. The key novel ingredient is that citizens are not purely selfish but rather, in line with the empirical evidence developed in Section 2, they have prosocial motives: they are averse to free riding on the costly voting efforts of others, inducing them to take actions that would be irrational from a purely selfish perspective. Specifically, we consider citizens who realize higher payoffs from voting when (i) they expect more supporters of their preferred candidate to vote, (ii) such votes collectively matter more for their preferred candidate's chance of winning, and (iii) they care more about which candidate wins. We show that each of the three components of prosocial motives modeled are needed to reconcile all identified empirical regularities.

We consider a canonical spatial setting with two candidates, where voting is costly. There is a continuum of two types of citizens, prosocial citizens and noise voters. The number of noise voters is random. A prosocial citizen's utility function can have two components: A private expressive (selfish) component from voting for their preferred candidate, and the prosocial component described above. In a Nash equilibrium, each prosocial citizen takes the voting behavior of others as given, and makes the turnout and voting decisions that maximizes their utility. We characterize equilibria that satisfy a weak refinement criterion that we term "no minimal expansion." This refinement rules out equilibria where an arbitrarily small set of abstainers could jointly switch to voting, making all of them strictly better off. This is a natural generalization of the feature of models with finitely-many citizens, where a citizen's decision can have at least a tiny impact on the probability of which electoral outcome obtains. The refinement ensures that the predictions of our continuum model correspond to those of models with a large, but finite, number of citizens.

We provide general conditions under which equilibria that satisfy no minimal expansion exist. We then show that a model specification that puts most of the weight on prosocial motives reconciles the following empirical facts that we document:

- High turnout even in lopsided elections. In our model, high turnout can obtain even with the continuum of agents, reflecting the presence of strategic complementarities in voting, as citizens are more willing to vote for their preferred candidates if others do the same.
- Higher turnout in closer elections. In our model, starting at a symmetric setting where each candidate is equally likely to win, shifting the distribution of citizen ideologies to favor, say candidate R discourages candidate L's supporters by more than it encourages R supporters to participate, lowering total turnout.
- Bandwagon effects (i.e., beliefs that a candidate is further ahead raise that candidate's support and reduce the opponent's). In our model, this again reflects strategic complementarities—if a citizen believes (e.g., due to favorable polls) that their candidate is likely to receive more support, this raises their expected prosocial payoff from voting, while the opposite happens for potential voters for the rival candidate.
- Voters with more extreme ideologies are less likely to abstain. In our model, turnout of voters with more extreme ideologies is higher because they care more about the electoral outcome than moderates. Using data from the American National Election Study (ANES), we establish semi-parametrically that voter abstention is single peaked in ideology. In particular, we find no evidence that more extreme voters turn out at lower rates.
- Increased voting costs or lower electoral stakes reduce turnout—this property is common to all models where citizens compare the costs and benefits of voting.

Importantly, we show that appropriately-parameterized versions of our model can *quantitatively* match voting patterns. We focus on off-cycle gubernatorial elections between 1998 and 2022 with no senate race, making the gubernatorial election the lead election and ensuring comparable electoral stakes. We use pre-election ratings from Cook Political Reports to establish four key features of toss-up and safe elections that we want our model to match: (i) turnout is substantial in both toss-up and safe races, but (ii) it is higher in the former—average voter participation falls from 49.3% in toss-up races to 41.2% in safe ones. This is so even though (iii) the average unsigned win margin rises from 5.3% in toss-up races to 29.5% in safe ones, and the Cook Political Reports correctly predict the outcomes of all safe elections, indicating that (iv) there is minimal uncertainty about who wins such elections. Our model can simultaneously match these four features using the same set of primitive parameters.

Each of the three components of prosocial preferences is needed for our model to quantitatively match established voting patterns. First, citizens must be more willing to participate when more like-minded citizens vote, even holding fixed the impact of those votes on the electoral outcome. This participation component helps sustain substantial turnout for a candidate who has little chance of winning. Second, the prosocial payoff from voting must also depend on the electoral effectiveness of co-supporters' collective voting efforts. This electoral-impact component generates a meaningful decline in total turnout as elections become less competitive, because the collective efforts of underdog supporters have less influence on their candidate's probability of winning. Finally, prosocial citizens must be more willing to vote when they care more about which candidate wins. Absent this third component, ideology would not affect turnout decisions and abstention rates would not be a single-peaked function of ideology.

We derive additional testable implications of our model for which we are not aware of causal empirical evidence. In particular, we predict that a symmetric increase in elite polarization increases turnout by raising the electoral stakes. Increasing mass polarization via a mean-preserving spread of voter ideologies has a similar effect. We further show that if, say candidate R's platform becomes more extreme, then this always raises participation of candidate L supporters, but can have an ambiguous effect on turnout for R.

Related Literature. Classical models of strategic voting consider fully rational and instrumental agents, who vote only if the probability that their vote is pivotal in determining the winner of the election is sufficiently high relative to the cost of voting (Riker and Ordeshook (1968); Palfrey and Rosenthal (1983, 1985); Myerson (1998); Castanheira (2003)). These pivotal voter models have a serious weakness—the vanishingly low turnout levels that they predict in large elections—the so-called *paradox of voting*. A common alternative, the expressive voting model, assumes that citizens enjoy the act of voting for their preferred candidate, which they compare to their cost of voting when deciding whether to participate. While research has shown that voters derive a direct "expressive" benefit from voting for a candidate whose platform closely aligns with their preferences,[1] this model's predictions are also at odds with facts established in Section 2.[2]

In reaction to the weaknesses of those widely-used models, researchers have proposed group-based voting models. There are three main variants: (i) ethical voter models assume that individuals are rule-utilitarians who decide whether to vote depending on what is optimal for their group as a whole (Feddersen and Sandroni,

[1]Empirical studies based on election data (e.g., Fujiwara (2011); Pons and Tricaud (2018); Spenkuch (2018)) and laboratory experiments (e.g., Forsythe et al. (1996); Esponda and Vespa (2014); Bouton et al. (2017)) provide strong evidence that voters receive a direct payoff from supporting a particular candidate while also caring about the election outcome.

[2]Further, in elections with more than two candidates, many voters cast misaligned votes—voting for a candidate other than their most-preferred one (contrary to the predictions of expressive voting models)—Abramson et al. (2018) finds that the incidence of voting for a second-choice candidate was almost 40% in some constituencies for the 2010 British general election, while Daoust (2018) finds that 22.6% of voters selected their second-choice in the 2015 Canadian general election.

2006; Coate and Conlin, 2004); (ii) mobilization voting models assume that individuals' turnout decisions depend on the mobilization efforts by leaders (Uhlaner, 1989; Morton, 1991; Shachar and Nalebuff, 1999); and (iii) social pressure models assume that individuals' voting decisions are influenced by pressure from peers (Ali and Lin, 2013; Levine and Mattozzi, 2020). Group-based models can deliver high turnout levels. However, we show these models fail to generate key patterns in the data, including bandwagon effects, large victory margins in races with high turnouts, and the relationship between ideology and abstention.

2 Empirical Regularities About Turnout

As the introduction observes, a canonical model of voting does not exist. A key issue is that existing models predict voting behavior that does not align with core features of turnout data. These empirical regularities concern the influence of various factors on turnout: voter ideology, voting costs, the expected closeness of an election, and whether the preferred candidate is the favorite or not. In this section, we review the existing evidence and provide novel evidence about the empirical relationship between citizen ideology and turnout/abstention, as well as patterns associated with toss-up versus safe electoral districts. We use these empirical regularities—both existing and novel—to assess the validity of our model's predictions, and to compare the performance of our model with those of other models.

2.1 Substantial Turnout Rates in both Competitive and Lopsided Elections

At the national level, the Global Dataset on Turnout (Coma and Van De Maele, 2023) considers elections for 190 countries for the 1945-2020 period. Across those elections, the mean turnout rate of registered voters is around 70% in both parliamentary and presidential elections. What is hard for many existing models to reconcile is that turnout remains high even if the outcome of the election is almost certain.

To illustrate how aggregate turnouts in very lopsided elections compare with those of closer elections, we consider non-presidential election cycles that featured only a gubernatorial race, making it the lead race to drive turnout, and ensuring that electoral stakes of races are comparable. Overall turnout data is taken from the University of Florida Election Lab.[3] Statewide election results are taken from `https://github.com/fivethirtyeight/election-results`. Ratings are taken from Cook Political Reports.

Figure 1 presents the relationships between the Cook reports rating of a race and realized turnout. We compare races categorized ex-ante by Cook reports as toss-ups with races considered solid for the Democrat or Republican candidate. In all races rated as solid, the favorite always won, indicating that they are extremely safe. There is a large variance in turnout for both toss-up and safe races, but even in safe races average turnout remained substantial at 41.2%, not much lower than the average of 49.3% in toss-ups.

[3] Michael McDonald. 2024. "1980–2022 General Election Turnout Rates (v1.1)." `https://election.lab.ufl.edu/dataset/1980-2022-general-election-turnout-rates-v1-1/`

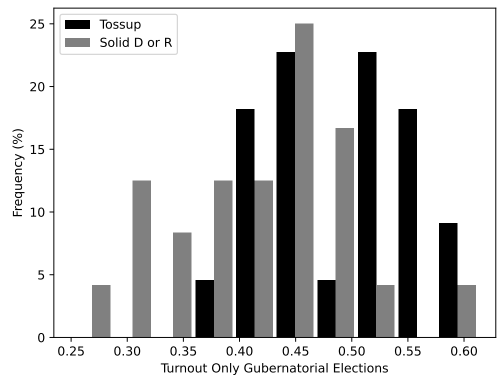

Figure 1: Turnout of eligible voters in races rated as toss-up vs races considered solid for one candidate, as determined by Cook Political Reports.

Even for candidates who have no reasonable chances of winning, voters still turn out in meaningful numbers. The left panel of Figure 2 (which shows the percentage of eligible voters who turn out and vote for losing candidates in elections considered as solid for the leading candidate) reveals that, on average, 15.5% of citizens still turn out to vote for candidates with no reasonable chance of winning—representing a substantial share of the total turnout in such elections. The right panel of Figure 2 shows that win margins for races rated as toss-ups are substantially smaller than those for races rated as solid. The associated average win margins are 5.3% versus 29.5%, where the win margins are calculated with respect to the runner-up candidate.

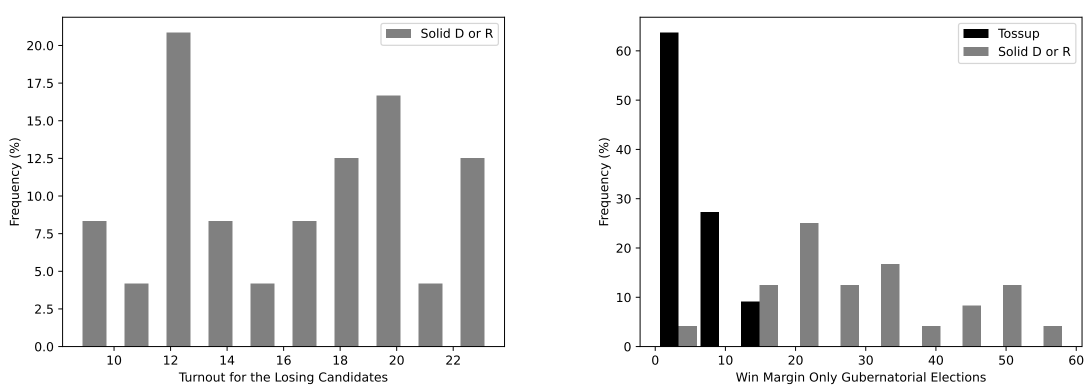

Figure 2: Turnout for losing candidates in races considered solid for one candidate (left), and the winning margin for races rated as toss-up vs races considered solid, as determined by Cook Political Reports (right).

2.2 Cost of Voting and Electoral Stakes

There is strong causal evidence that reducing voting costs or raising electoral stakes increases turnout. Turnout is much higher in presidential election years in the US than in off-cycle election years, illustrating the impact of higher electoral stakes. To quantify the impact of voting costs, some studies explore the impact of voting laws about early voting (Kaplan and Yuan, 2020) and voting-by-mail (Thompson et al., 2020; Yoder et al., 2021) on turnout. These U.S. studies exploit plausibly exogenous variations in voting laws across counties or across individuals in a given constituency to identify a causal impact of the laws on turnout. They find that turnout increases when voting laws make the act of voting simpler, and hence less costly. In field experiments, Braconnier et al. (2017) and Harris et al. (2025) show that making voter registration easier raises turnout. Other studies find that rain on election day—which increases voting costs—reduces turnout (Gomez et al., 2007; Concha-Arriagada and Naddeo, 2022). Finally, studies use border discontinuity designs to identify the causal impact of distance to polling station on turnout, finding that turnout decreases with the distance to polling stations (Cantoni, 2020; Tomkins et al., 2023; Bagwe et al., 2022).[4]

2.3 Closeness of Elections

The empirical literature identifies a strong positive correlation between the expected closeness of elections and turnout. For instance, Blais (2000)'s review found that in 27 of 32 studies, turnout was higher in elections that were expected to be closer. Still, establishing the causal impact of closeness requires more, as the positive correlations could stem, e.g., from systematic differences in voter characteristics in countries or electoral districts with close vs. lopsided elections.

More recent studies use a variety of strategies to estimate the causal impact of expected closeness on turnout. In US elections, Moskowitz and Schneer (2019) and Ainsworth et al. (2024) exploit changes in the closeness of congressional electoral districts due to redistricting. Using a nationwide panel of voters, Moskowitz and Schneer (2019) report a null effect of congressional district closeness on voter turnout. Ainsworth et al. (2024) focus on North Carolina, and find a small positive effect. In US presidential races, Miller (2025) provides quasi-experimental evidence on the effect of competitiveness on turnout, comparing adjacent counties in different states but within the same media market (i.e., fixing campaign advertising exposure), leveraging sharp discontinuities in statewide competitiveness. He finds that moving from average to high levels of competitiveness raises turnout by 1.67 percentage points. For Swiss elections, Bursztyn et al. (2024) use an event study design that exploits changes in polling data and precise information about the day mail ballots being cast. They also find a small positive effect of expected closeness on turnout.

Other researchers use field and laboratory experiments to estimate the causal impact of election closeness on turnout. For laboratory experiments, Agranov et al. (2018) summarizes the literature as follows: "experimental work has suggested that higher probabilities of pivotality indeed induce greater participation rates." In

[4] Other studies of the impact of distance to polling station include Brady and McNulty (2011); Yoder (2018); Clinton et al. (2021).

other words, turnout increases with closeness. For field experiments, Enos and Fowler (2014); Gerber et al. (2020) detect slight responses of turnout in the US to random provision of information on the closeness of the race.

Overall, these empirical findings suggest a small positive causal impact of closeness on turnout.

2.4 Underdog and Bandwagon Effects

Underdog and Bandwagon effects refer to changes in the voting propensity of a given individual caused by changes in beliefs about a candidate's winning probability. A bandwagon effect arises if the voting propensity of a given individual *increases* with their beliefs that their favorite candidate will win. An underdog effect arises if their voting propensity *decreases* with their beliefs that their favorite candidate will win.

The experimental literature presents compelling evidence of bandwagon effects (Duffy and Tavits, 2008; Großer and Schram, 2010; Kartal, 2015; Agranov et al., 2018). There are three types of evidence. First, some studies compare the voting propensity of supporters of the (expected) frontrunner and the (expected) runner up. These studies find that supporters of the frontrunner have a higher voting propensity than supporters of the underdog. Second, some studies assess the impact of revealing information about the distribution of preferences on voting propensities. They find a positive impact on the voting propensity of supporters of the candidate revealed to be the frontrunner, and a negative impact on supporters of the revealed runner-up. Finally, other studies elicit subjects' beliefs about the probability of winning of the two candidates. They find that the voting propensity increases with the subjects' beliefs about their preferred candidate's electoral advantage.

Outside of the laboratory, the challenge is to identify shocks on beliefs about candidates' winning probabilities, and how they impact voting propensity of a given set of individuals. To overcome this challenge, Fraga et al. (2022) combine redistricting episodes with individual-level data on turnout decisions. The redistricting episodes act as shocks on the partisan composition of Congressional districts (and hence on candidate winning probabilities). The individual-level data let them explore the change in behavior of the individuals over time (using individual fixed effects). They find evidence of bandwagon effects: increasing the partisan lean of congressional districts increases the voting propensity of the supporters of the favored candidate, even if that candidate has an overwhelming advantage due to gerrymandering.[5] Of note, this implies that the endogenous participation responses of voters accentuates the impact of a candidate's baseline support advantage on their chances of winning and actual expected win margin.

[5]Lehmann and Lindlacher (2026) exploit the run-off structure of elections in Germany to provide additional causal evidence of bandwagon effects. They find that if, in the first round, the AfD gained more votes then in the runoff election the AfD gains more votes, while the rival gains less, controlling for extremism of the rival.

2.5 Citizen Ideology

Different models of turnout make distinct predictions about the relationship between ideology and turnout. For example, purely expressive voting predicts a W-shaped abstention pattern: turnout is highest for voters whose ideal points are close to their preferred party's platform, and lower when they are further away. By contrast both our model and the pivotal voter model predict a V-shaped turnout pattern: turnout is highest for more extreme voters, and lowest for more moderate ones.

These predictions can only be tested if one has a sufficiently fine measure of ideology, so that non-linear patterns can be observed, even if they only occur for some of the more extreme types. Existing empirical research on the relationship between ideology and turnout takes one of two approaches that preclude distinguishing between the predictions of different models. For example, Ortoleva and Snowberg (2015) estimate a *linear* relationship between the degree of extremism in ideology and turnout, thereby precluding obtaining a W-shaped pattern.[6] Alternatively, researchers impose functional relations between ideology and turnout that preclude generating a W-shaped abstention pattern, for example, by structurally estimating a model in which citizens only vote if their utility difference between candidates is sufficiently high (e.g., Kawai et al. (2021)[7]).

We adopt a semi-parametric approach that lets the data reveal how the actual turnout pattern varies with ideology. The only structural assumption that we make is that voters have quadratic preferences and are subject to normally-distributed valence shocks. Importantly, we do not impose a model of turnout. Instead, we first use information about citizens who vote for one of the two main candidates to infer their ideologies. We then infer the ideology of citizens who do not vote from the parameter estimates of voting citizens. Our approach only implicitly assumes that (i) voters and non-voters with identical observable characteristics have the same ideology, and (ii) if heterogeneity in voting costs matters for turnout, then while voting costs can depend on ideology, the converse is not true—voting costs do not causally impact voter ideology. We now detail our estimation approach.

Our analysis uses survey data from the American National Election Studies (ANES). Key features of these data are that a large common set of questions about political preferences is asked over multiple years, information about turnout is validated, and both sample size and response rates are large. We start by estimating the ideology of voters in U.S. presidential elections. For citizens who turn out and vote for one of the two main candidates, D and R, with associated platforms x_D and x_R, we assume that their voting choices are based on quadratic preferences around their respective ideal points and idiosyncratic valence shocks. Given that a citizen θ finds it optimal to vote, θ votes for candidate i over candidate j if

$$-(\theta - x_i)^2 + \epsilon_i > -(\theta - x_j)^2 + \epsilon_j,$$

where the ϵ_i and ϵ_j are independent, idiosyncratic, normally-distributed preference shocks for candidates i and

[6] Ortoleva and Snowberg (2015) do this because their main objective is to determine how overconfidence affects turnout.

[7] They do this because the focus of their estimation is to understand voter perceptions of probabilities of being decisive.

j. Thus, voter θ votes for candidate R when

$$\epsilon_D - \epsilon_R < 2(x_R - x_D)\theta - x_R^2 + x_D^2, \tag{1}$$

and θ votes for candidate D when this condition is reversed. In particular, *conditional on voting*, the further θ is to the left, the greater is the probability that θ votes for D.

To relate the ideology of voters to their answers to ANES survey questions (details below), we assume that $\theta = X \cdot \zeta + b + \delta$, where X is a vector of the respondent's answers, b is an intercept, and δ is drawn from a normal distribution. Substituting $X \cdot \zeta + b + \delta$ for θ in (1) and rearranging terms yields that voter θ votes for R if

$$\epsilon_D - \epsilon_R - 2(x_R - x_D)\delta < 2(x_R - x_D)X \cdot \zeta - x_R^2 + x_D^2 + 2(x_R - x_D)b. \tag{2}$$

As $\epsilon = \epsilon_D - \epsilon_R - 2(x_R - x_D)\delta$ is normally distributed, if we subtract $E[\epsilon]$ from both sides of (2), the random variable $\tilde{\epsilon} = \epsilon - E[\epsilon]$ on the left-hand side has mean zero. Defining $\tilde{b} = -x_R^2 + x_D^2 + 2(x_R - x_D)b - E[\epsilon]$, equation (2) takes the form

$$\tilde{\epsilon} < 2(x_R - x_D)X \cdot \zeta + \tilde{b}. \tag{3}$$

We thus have a probit model that we can estimate via maximum likelihood. The model identifies ideologies only up to an affine linear transformation. As a result, in each election year, we normalize the most liberal voter at -1 and the most conservative at $+1$. This normalization does not let us compare ideologies over time, but this is not necessary given our focus on the link between ideology and turnout rates for individuals in a *given* election.

To estimate the coefficients ζ, we need to specify which answers to ANES questions to include in the X vector. We use the same covariates in all election years to preserve comparability. We use as covariates the thermometer scores that measure a respondent's feelings toward various groups from 0 to 100, where 0 means feeling very "cold" and 100 means feeling very "warm". The thermometer scores were first available for the 2004 election. We initially used all of the thermometer scores in ANES as covariates but we then removed those that were not significant in 2020. We do this because some individuals do not answer all questions, and reducing the number of questions increases the sample size, which matters especially in earlier election years. In particular, while there are 4,900 observations of votes for the two major candidates in 2020 by individuals who answered all the thermometer questions, there are only 416 such observations in 2004. The thermometer scores included are attitudes towards: (i) liberals, (ii) conservatives, (iii) Democratic party, (iv) Republican party, (v) illegal aliens, (vi) Christian fundamentalists, and (vii) feminists.[8] We also include the self-identification on the liberal-conservative scale, which goes from 1 to 7.[9]

Table 1 presents the results from estimating the weighted probit regression for 2020. Estimated parameters for other election years are qualitatively similar. The model predicts voting choices extremely well, as

[8]Results are qualitatively identical if we include the non-significant variables, gender, attitudes towards Blacks, big business, unions, gays, and SCOTUS.

[9]We removed observations of "don't know" or "haven't thought much about it". Recoding these respondents as moderates gives rise to qualitatively similar estimates but substantially worsens the model fit, indicating that such recoding is inappropriate.

highlighted by the pseudo R^2 of 0.946 for the 2020 election sample. Adding the thermometer covariates to the liberal-conservative self-identification matters. Including them notably improves the model fit, e.g., raising the pseudo R^2 from 0.411 to 0.733 for 2004 and from 0.658 to 0.946 for 2020. The thermometer covariates also let us estimate a continuous distribution over θ, rather than a distribution with only 7 different types.

Variable	coef	std err	z	$P > \lvert z \rvert$
Intercept	0.4979	0.025	19.572	0.000
Feelings towards liberals	-0.0016	0.000	-7.229	0.000
Feelings towards conservatives	0.0010	0.000	4.652	0.000
Feelings towards Democratic party	-0.0051	0.000	-26.413	0.000
Feelings towards Republican party	0.0049	0.000	24.184	0.000
Feelings towards illegal aliens	-0.0012	0.000	-7.334	0.000
Feelings towards Christian fundamentalists	0.0006	0.000	3.802	0.000
Feelings towards feminists	-0.0010	0.000	-5.295	0.000
Liberal-Conservative self-identification	0.0148	0.004	3.765	0.000

Table 1: Results of Weighted Probit Regression for 2020. $N = 4,900$, Pseudo R^2 is 0.9459. All variables except for the liberal-conservative scales are thermometer scores on a 0–100 scale. The liberal-conservative self-identification values range from 1 (extremely liberal) to 7 (extremely conservative).

We then use the parameter estimates for voters to estimate the ideologies of non-voters (abstainers) by plugging in their thermometer scores into the estimated equation. In the final step, we estimate separately the density functions f_V and f_A of the ideology distributions of voters for the two major candidates and abstainers, respectively, using kernel density estimation (`gaussian_kde` from scipy.stats). The abstention probability for a type θ is given by $f_A(\theta)N_A/(f_A(\theta)N_A + f_V(\theta)N_V)$, where N_V and N_A are the sums of the sample weights of voters and abstainers in the survey, respectively.

Figure 3 plots how abstention rates vary with ideology for the election years 2004 through 2020.[10] The results for 2008 to 2020 reveal that abstention rates for centrists range from three to five times those for extremists. Moreover, there is no evidence that abstention rates rise in the tails. In 2004, the results are very noisy due to the modest number of observations (there are 60% fewer observations of voters in 2004 than in 2008 and over 90% fewer than in 2020) manifested in the especially-wide confidence bands in the tails, where there are very few observations. In sum, our findings comprise strong evidence that abstention occurs primarily among centrists, and we do not find evidence for increased abstention rates among extremists.[11]

[10] We bootstrap 95% confidence bands, repeating the estimations of ideologies and the subsequent kernel density estimations using resampling from the original data. For each ideology θ we plot the upper and lower 2.5 percentiles, given by the dotted lines.

[11] Abstention rates among those who complete the ANES are lower than in the election, reflecting the selection of motivated respondents to surveys. Before we apply filters for complete responses, ANES abstention rates range between 41% (in 2020) and 63% (in 2008) of those in the election. After filtering for complete answers to the relevant thermometer questions, abstention

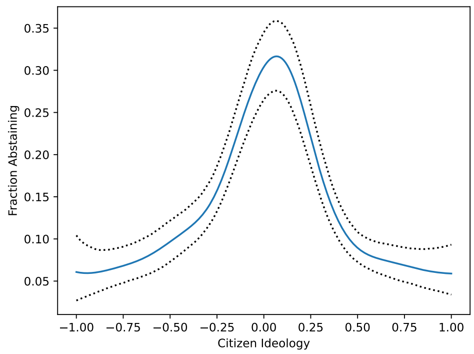

(a) 2020

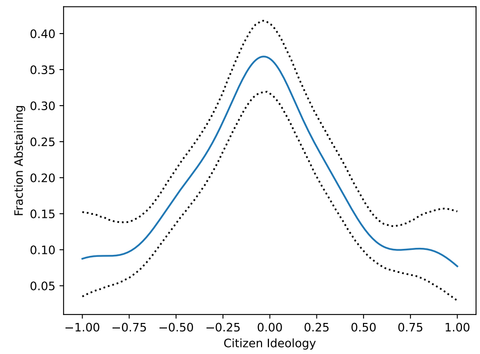

(b) 2016

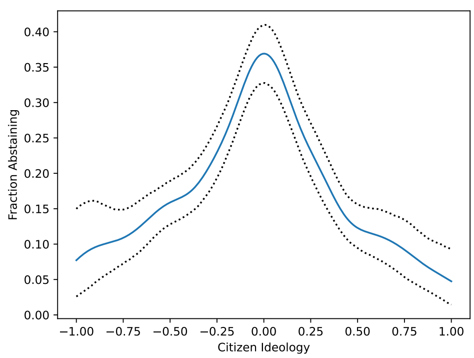

(c) 2012

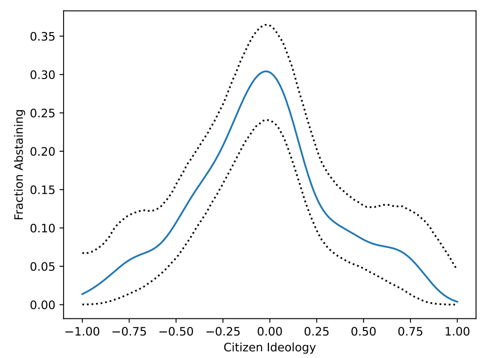

(d) 2008

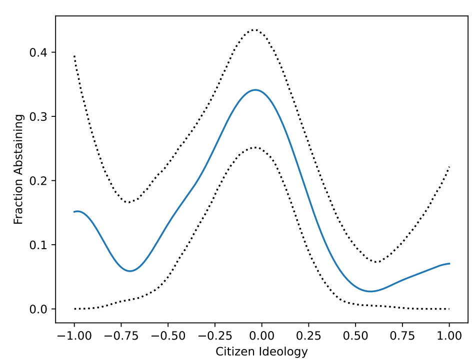

(e) 2004

Figure 3: Abstention rates of citizens by ideological positions with 95% confidence intervals.

rates range between 33% (in 2020) and 48% (in 2008) of the actual rates. The response rate to the 2020 survey was 40.8% (`https://electionstudies.org/wp-content/uploads/2022/08/anes_timeseries_2020_methodology_report.pdf`) down from 60.5% in 2000 (`https://electionstudies.org/data-quality/`). It is not surprising that individuals who respond to political surveys are more likely to vote. Further, respondents to political surveys likely have stronger political opinions, in which

3 The Model

There are two candidates, $i = L, R$, with policy positions $x_L < x_R$, and a continuum of citizens. There are two types of citizens—prosocial citizens and noise voters. A noise voter for candidate i always votes for i regardless of the policy positions taken. The number of noise voters for each candidate is random. Noise voter support for a candidate i is summarized by the number $\rho_i \geq 0$ of i's noise voters. We define $\rho = \rho_R - \rho_L$ to be the net noise voter turnout for candidate R, which we assume has distribution G.

Prosocial citizens are distinguished by their ideal policy bliss points, y, which we refer to as their ideologies. Let Φ be the cdf over y in $\mathbb{R}$. With slight abuse of notation we also use Φ to denote the probability measure. We assume that Φ has a density ϕ that is continuous and strictly positive on its connected support.

Each prosocial citizen has three possible strategies: vote for L, vote for R, and abstain. A strategy profile is thus fully described by the sets of prosocial citizens S_L and S_R who vote for candidates L and R, respectively. Formally, $S_L, S_R \subset \mathbb{R}$, with $S_L \cap S_R = \emptyset$; the set of abstainers is $\mathbb{R} \setminus (S_L \cup S_R)$. Thus, candidate L wins with probability $P(S_L, S_R) = G(\Phi(S_L) - \Phi(S_R))$, while candidate R's winning probability is $1 - P(S_L, S_R)$.

The utility that a prosocial citizen with policy bliss point y from voting for candidate i is given by

$$u_y(i, S_L, S_R; x_L, x_R) = \beta v_y(i; x_i) + (1 - \beta) B_y(i, S_L, S_R; x_L, x_R) - c.$$

Here $v_y(i; x_i)$ is the standard spatial utility payoff that y derives from voting when i has policy platform x_i; B_y is the prosocial benefit that y derives from voting for i; and $c > 0$ is the cost of voting.[12] Parameter β controls the level of prosocial considerations, with $\beta = 1$ capturing purely selfish preferences and $\beta = 0$ capturing purely prosocial preferences. The key feature of prosocial preferences is that individuals get more utility from prosocial behavior if more co-supporters participate, i.e., B_y is monotone in S_i:

Definition 1 *B_y is monotone if, for citizens y who prefer candidate i the following holds: $B_y(i, S'_i, S_{-i}; x_L, x_R) \geq B_y(i, S_i, S_{-i}; x_L, x_R)$ for $S'_i \supset S_i$. B_y is strictly monotone if the inequality is strict for $\Phi(S'_i) > \Phi(S_i)$.*

Thus, a citizen is more willing to vote for their preferred candidate, i, instead of abstaining, if they expect more other citizens to incur the cost of participating by also voting for i. We assume that the utility of a prosocial citizen y who abstains is independent of y and c, and hence can be written as $u(N)$. It is routine to extend the analysis to allow the utility of abstaining to vary with y. In what follows, we drop the dependence of u_y and B_y on x_L and x_R in our notation where it does not cause confusion.

Finally, we need to ensure that utility is consistent with the spatial structure of our model:

Assumption 1 *$u_y(R, S_L, S_R; x_L, x_R) - u_y(L, S_L, S_R; x_L, x_R)$ is strictly increasing in y.*

case the true abstention pattern would be an even steeper single-peaked function of ideology than what we find.

[12]In Section 7 we show that the model can be easily extended to allow for heterogeneous voting costs.

Of course, this assumption is satisfied by standard quadratic preferences, but it is significantly more general.

A Nash equilibrium in our game is a strategy profile (S_L, S_R) such that every prosocial citizen y is best responding to the behavior of other prosocial citizens:

Definition 2 *(S_L, S_R) is an equilibrium if and only if all prosocial citizens maximize their utility given S_L and S_R, i.e., for all $y \in S_i$, $u_y(i, S_L, S_R; x_L, x_R) \geq u(N)$ and $u_y(i, S_L, S_R; x_L, x_R) \geq u_y(-i, S_L, S_R; x_L, x_R)$. If $y \notin S_L \cup S_R$ then $u(N) \geq u_y(i, S_L, S_R; x_L, x_R)$, for $i = L, R$.*

To ease analysis, we focus on a model with a continuum of prosocial citizens. However, we want to be able to interpret our equilibrium as the limit of a polity with a large but finite number of prosocial citizens. In a finite setting, any individual's decision to vote for a particular candidate strictly increases turnout for that candidate. This latter effect increases the citizen's willingness to vote in a finite model, but it is not present in a continuum model. In our setting this can lead to Nash equilibria that are not robust in the sense that even an infinitesimally small set of voters could improve by jointly deviating from non-voting to voting.

To address this, we focus on equilibria that satisfy the following property:

Definition 3 (No Minimal Expansion) *Let (S_L, S_R) be an equilibrium. No minimal expansion holds if expanding S_i marginally does not make all new voters strictly better off: For all $\bar{y}_i$ on the boundary of S_i,*
if there exists $y > \bar{y}_i$ with $(\bar{y}_i, y) \cap (S_L \cup S_R) = \emptyset$ then

$$\frac{\partial u_y(i, S_i \cup [\bar{y}_i, y], S_{-i}; x_L, x_R)}{\partial y}\bigg|_{y=\bar{y}_i} \leq 0,$$

if there exists $y < \bar{y}_i$ with $(y, \bar{y}_i) \cap (S_L \cup S_R) = \emptyset$, then

$$\frac{\partial u_y(i, [y, \bar{y}_i] \cup S_i, S_{-i}; x_L, x_R)}{\partial y}\bigg|_{y=\bar{y}_i} \geq 0.$$

To see how this definition corresponds to the intuitive formulation that expanding the set of voters by an infinitesimal amount does not make the new voters better off, suppose that S_L is expanded from $(-\infty, y_L)$ to $\tilde{S}_L = (-\infty, y_L + dy_L)$ and the new members are strictly better off, i.e., $u_{y_L+dy_L}(L, \tilde{S}_L, S_R; x_L, x_R) - u(N) > 0$. Dividing by dy_L and using the fact that in equilibrium the marginal voter in S_L is indifferent between voting and not voting, we get $(u_{y_L+dy_L}(L, \tilde{S}_L, S_R; x_L, x_R) - u_{y_L}(L, S_L, S_R; x_L, x_R))/dy_L > 0$, which violates the definition.

In addition, we need B_y to be concave or strictly concave in y.

4 Discussion of Modelling Assumptions

4.1 Prosocial Motives and Voting

There are many real-life examples of prosocial behavior.[13] Beyond these examples, a large body of evidence shows that people cooperate far more than what standard economic theory predicts. For example, an extensive experimental literature finds that many individuals do not free-ride even when, from a rational and selfish perspective, it is a dominant strategy to do so (see the reviews by Gachter and Herrmann (2009); Fehr and Schurtenberger (2018)).[14]

Malmendier et al. (2014) classify theories of prosocial behavior according to the "internal" or "external" nature of the factors underlying the prosocial behavior. Internal factors relate to individual preferences and encompass altruism, warm-glow, ethical concerns, and inequity aversion, as well as the reciprocity they may generate (Andreoni and Miller (2002); Charness and Rabin (2002); Falk and Fischbacher (2006); Fehr and Schmidt (1999); Bolton and Ockenfels (2000); Levine (1998); Rotemberg (2008); Battigalli and Dufwenberg (2009); Rabin (1993)). External factors relate to social considerations such as social pressure and punishment, social image and norms, including reciprocity norms (Bénabou and Tirole (2006); Battigalli and Dufwenberg (2007); Andreoni and Bernheim (2009); Ali and Lin (2013); Fehr and Gächter (2002)).[15]

The literature finds that both internal and external factors motivate citizens to behave pro-socially when deciding whether to vote. For internal factors, various studies use survey data to explore voter motives (Blais (2000); Knack (1992); Opp (2001); Jankowski (2007)), highlighting altruism and civic duty as key motives. Using laboratory experiments, Fowler (2006) shows that altruism plays a key role in turnout decisions. For external factors, Bond et al. (2012) and Gerber et al. (2008, 2016, 2017) use field and survey experiments to provide clear evidence of the relevance of social considerations for voting. They highlight that participation increases substantially following information about the behavior of one's neighbors or promises to share information about one's turnout with neighbors. DellaVigna et al. (2017) use a field experiment to provide direct evidence that citizens vote to tell others (or not to be forced either to admit that they abstained, or to lie about it), and more generally that image concerns matter for voting behavior.[16] Funk (2010) uses a natural experiment in Switzerland after the introduction of optional postal voting to provide evidence about the relevance of social motives. She finds that the impact of introducing optional postal voting is smaller (either less

[13] For instance, millions routinely donate blood, incurring discomfort with no monetary reward; people often recycle, reduce water usage, or pay premiums for sustainable products, even when the personal benefits are tiny and the actions only cumulatively matter if many others also cooperate; in most countries, a large fraction of the electorate votes even though the probability that an individual ballot changes the electoral outcome is essentially zero; many people report income and pay taxes even when the risk of audit is tiny.

[14] These results are established in public good experiments where players simultaneously decide how much of their private endowment to contribute to a public good and how much to keep for themselves, creating a social dilemma between the individually-rational incentive to free-ride (contribute nothing) and the collectively-optimal outcome of full contribution by everyone.

[15] For experimental evidence of the importance of social considerations for prosocial behavior, see Fehr and Gächter (2000); Fehr and Gächter (2002); Cubitt et al. (2011); Kimbrough and Vostroknutov (2016), and references therein.

[16] Silver et al. (1986) and Belli et al. (1999) also find that a substantial share of abstainers lie about having abstained when asked.

positive, or even negative in cantons with legal duty to vote) in small communities where social incentives are stronger than in large communities (where social incentives are arguably weaker).

4.2 Monotonicity of B_y

We model prosocial motives in a reduced-form way, directly including them in the B_y function. This function modulates the benefits to prosocial citizen y from voting relative to abstaining as a function of the behavior of other citizens. A key feature of B_y is that it is monotone in S_i for i voters: prosocial motives depend on how many fellow citizens vote for one's preferred candidate. In particular, they are more prosocial/cooperative if they believe that others will also be prosocial/cooperative, and they are more selfish if they believe that others will be selfish. Citizens thus have "conditional prosocial motives".[17] This key feature is in line with the idea of "conditional cooperation," which originates in the social psychology literature (see Kelley and Stahelski (1970)) and has been documented extensively in the experimental literature. Using the "strategy method," this literature provides evidence that many subjects are conditional cooperators, as they cooperate at high levels if they believe others are doing the same, but free ride if they believe others are not cooperating (Fischbacher et al. (2001), Kocher et al. (2008) and Chaudhuri (2011)). While most of the experimental literature focuses on public good provision games, researchers also document evidence of the relevance of conditional cooperation in "group contest" settings, which encompass voting situations (e.g., Kiss et al. (2020)).

The literature provides direct evidence of the conditionality of prosocial motives in voting. Gerber and Rogers (2009) use get-out-the-vote field experiments to show that messages emphasizing high expected turnout by other citizens are more effective at increasing turnout than messages emphasizing low expected turnout. These findings align with the core idea of conditional cooperation that people's desire to cooperate depends on what they believe others will do. Doherty et al. (2017) use survey experiments to explore how the expected behavior of others affects one's decision to turnout. They document a positive relationship between the perceived social benefit of voting and the expected voting behavior of others.

4.3 Types of Prosocial Motives Captured by B_y

Our specification of a monotone B_y function is general enough to encompass several distinct motives underlying prosocial voting. This flexibility is important because, as discussed above, the empirical and experimental

[17]Many of the external and internal factors underlying prosocial motives in voting give rise to conditional prosocial motives. For instance, if social pressure or punishment for abstention comes primarily from co-supporters of a candidate, then the cost of abstention rises with the number of co-supporters who vote. The same holds for social image concerns if their intensity is shaped by citizens voting for one's favorite candidate. On the internal factor side, altruistic citizens may see abstention as less appealing when more co-supporters vote, as voting becomes a way to compensate them for the costs they incurred. Similarly, citizens motivated by warm-glow from "doing their part" may perceive voting as more appealing when more co-supporters do their part by voting. Quite broadly, reciprocity—whether due to internal or external factors—makes abstention less appealing relative to voting when more co-supporters vote.

literatures suggest turnout may be driven by multiple motives, including altruism, civic duty, warm-glow, and reciprocity, as well as external motives, such as social pressure or image concerns, and norms of reciprocity. Our approach does not require taking a stand on which microfoundations matter most. Instead, B_y captures the basic feature that a citizen's benefit from voting for her preferred candidate may depend on the costly voting efforts of like-minded citizens, while allowing the drivers to take different forms.

First, B_y can capture prosocial benefits from voting that depend on the number of co-supporters who vote, independently of whether their votes are likely to change the electoral outcome. This includes warm-glow from "doing one's part," civic-duty motives that are stronger when others also participate, and reciprocity toward fellow supporters who incur the cost of voting. It also encompasses social-image and social-pressure motives when the relevant audience or reference group consists primarily of citizens who support the same candidate. In all of these cases, the benefit from voting rises with $\Phi(S_i)$, holding fixed the marginal effect of these votes on candidate i's winning probability.

Second, B_y can capture prosocial benefits from voting that depend on the extent to which the costly voting efforts of co-supporters collectively affect the preferred candidate's probability of winning. A citizen may dislike free-riding on the efforts of like-minded citizens when those efforts matter more for the electoral outcome. As a result, the benefit from voting is higher when the participation of co-supporters of candidate i has a larger effect on i's winning probability. This motive links prosocial voting to the electoral effectiveness of the group effort.

Third, B_y should capture prosocial benefits from voting that depend on how much the citizen cares about which candidate wins. A citizen's willingness to incur the cost of voting for others should be stronger when the utility difference between the two candidate platforms, $v_y(i) - v_y(-i)$, is larger. This component links prosocial behavior to the electoral stakes faced by citizen y, and implies that citizens with stronger policy preferences may derive larger prosocial benefits from supporting their preferred candidate.

We next establish the existence of an equilibrium. The proof of existence is general in nature and does not necessitate in any way that B_y incorporates any of the three specific features highlighted above.

5 Existence

To prove existence of equilibria that satisfy no minimal expansion, we require continuous differentiability of utility for all cases where the S_i are intervals.

Assumption 2 *Let* $D = \{(y, y_{1L}, y_{2L}, y_{1R}, y_{2R}) \in \bar{\mathbb{R}}^5 | y_{1j} \leq y_{2j},\ j = L, R, y_{2L} \leq y_{1R}\}$. *Then the function* $(y, y_{1L}, y_{2L}, y_{1R}, y_{2R}) \mapsto u_y(i, [y_{1L}, y_{2L}], [y_{1R}, y_{2R}])$ *is continuously differentiable on D.*

Note that continuous differentiability must hold at $\pm\infty$, and also at $y_{iL} = y_{iR}$, in which case the one-sided derivative is taken. For example, if utility only depends on the measure of S_i, which we can denote by m_i, then the assumptions are satisfied if $u_y(i, m_L, m_R)$ is continuously differentiable on $\mathbb{R} \times [0, 1]^2$ and the density ϕ of the distribution of prosocial citizens is continuous.

Note that in a continuum model, the decision of a single citizen whether or not to participate should not affect the payoff of the other citizens. This is obviously true if utility takes the form $u_y(i, m_L, m_R)$. However, for more general specifications of utility we assume that this property holds.

Assumption 3 *A citizen's utility does not vary with the participation decision of any single other citizen: $u_y(j, S_i, S_{-i}) = u_y(j, S_i \cup \{y'\}, S_{-i})$, for all citizens $y \neq y'$ and S_i, $i = L, R$.*

Using the above assumptions we have:

Theorem 1 *Let v_y and B_y be concave in y. Further, suppose that B_y is either strictly monotone, or monotone and strictly concave in y. Then there exists a pure strategy Nash equilibrium in which S_L and S_R are either intervals or empty, and for which no minimal expansion holds.*

Conversely, in any Nash equilibrium, the set of citizens S_i who vote for candidate i with strictly positive probability is an interval or empty, and all citizens in the interior of S_i must vote for candidate i with probability 1.

To prove Theorem 1, we approximate the continuum economy by a sequence of economies using an interval partition of the support of Φ, where each interval has the same measure. The midpoint of each interval is the representative voter for that interval, and thus has positive mass. Each of the approximating economies has a mixed strategy Nash equilibrium, according to Nash's existence theorem. We then show that the concavity and monotonicity assumptions imply that the set of voters is a "discrete" interval (no holes), and that mixing can only occur at the interval boundaries. Standard compactness arguments imply that a subsequence of these intervals converge, and an upper hemicontinuity argument implies that the limit is an equilibrium of the continuum economy. This argument uses the continuity of utility when the S_i take the form of intervals, which follows from Assumption 2. Using Assumption 3 we can purify this limit equilibrium. In particular, citizens at the boundary of S_i must be indifferent between voting and abstaining. We can assume that they vote with probability 1 as they are indifferent, and the payoff of other citizens are unaffected by Assumption 3. Finally, the property of no minimal expansion is inherited from the finite approximating economies.

6 Comparative Static Analysis

6.1 Parametric Specification of B_y

In the remainder of the paper we consider a simple parametric specification which contains the three features of prosocial voting preferences highlighted in Section 4.3. We will show that to explain the data qualitatively, monotonicity of B_y in S_i is sufficient. To match the quantitative features of the data, we need all three features.

Let

$$B_y(i, S_L, S_R; x_L, x_R) = V\Big[(1-\lambda)\Phi(S_i) + \lambda\Delta_i(S_i, S_{-i}; y)\Big](v_y(i) - v_y(-i)). \tag{4}$$

Here $0 \leq \lambda \leq 1$ governs the relative weight placed on the electoral-impact motive rather than the participation motive, $V > 0$ is the stake of the election, and $\Delta_i(S_i, S_{-i}; y)$ measures the change in candidate i's winning probability generated by the participation of the relevant set of co-supporters from y's perspective. In particular, one can write

$$\Delta_L(S_L, S_R; y) = P(S_L, S_R) - P(S'_L(S_L, y), S_R),$$

and

$$\Delta_R(S_R, S_L; y) = (1 - P(S_L, S_R)) - (1 - P(S_L, S'_R(S_R, y))),$$

where $S'_L(S_L, y)$ are the co-supporters of L that y does not care about, and $S'_R(S_R, y)$ is defined similarly.

Thus, $\Phi(S_i)$ captures prosocial benefits that depend directly on the number of co-supporters who vote; $\Delta_i(S_i, S_{-i}; y)$ captures prosocial benefits to y that depend on the electoral impact of the voting efforts of the co-supporters that y cares about, $S_i \setminus S'_i(S_i, y)$; and $v_y(i) - v_y(-i)$ captures prosocial benefits that depend on the intensity of citizen y's preference over the electoral outcome.

There are many possible formulations of $S'_i(S_i, y)$. Monotonicity requires that expanding the set of citizens voting for candidate i weakly increases the prosocial benefit from voting for i. Concavity in y underlies the existence result ensuring that equilibrium participation sets have an interval structure. This restriction is trivially satisfied when $\Phi(S'_i)$ is a fixed share of $\Phi(S_i)$. By contrast, concavity can fail if citizen y cares only about the participation of other citizens within a fixed ideological distance. With a strictly single-peaked distribution of types, centrist citizens have more such neighbors than more extreme citizens, which can generate non-concavities and non-interval structures for the set of voters, which we do not see in the data. In our comparative static analyses, we adopt the simplest formulation, $S'_i(S_i, y) = \emptyset$, which satisfies these requirements and is sufficient to match the data qualitatively and quantitatively.

6.2 Qualitative Analysis

6.2.1 Parametric Specification

We next show analytically that the simplest specification of B_y in (4) , with $\beta = \lambda = 0$ and v_y quadratic, already qualitatively delivers all of the empirical regularities identified in Section 2.[18] This reflects that most of the arguments follow from monotonicity of B_y in S_i, which, as discussed, reflects the conditional cooperation embedded in prosocial behavior. More generally, qualitative matching requires that β be sufficiently small, i.e., prosocial motives are the main driver of participation. The results are robust to $\beta > 0$ and we show that

[18]Our analytical characterizations extend immediately to $\lambda > 0$ when G is uniform.

increasing β reduces turnout. In Section 6.3 we show that $\lambda \gg 0$ is necessary to quantitatively match the data, i.e., citizens must care significantly about the impacts of like-minded citizens on the electoral outcome.

To ease analysis, we make the following assumptions:

Assumption 4 *The distribution Φ of ideologies has a single-peaked and log-concave density, ϕ.*

Consider a citizen, y, who is located at the midpoint between the candidate positions. This citizen's utility of voting is $-c$. The utility of abstention is u_N. If $-c > u(N)$ then citizen y and hence all other citizens participate. As 100% turnout is not a realistic scenario, we assume in the following that $-c < u(N)$.

6.2.2 Voter Participation

Theorem 1 shows that if (S_L, S_R) is an equilibrium, then the sets S_i are either intervals or empty. When $\beta = 0$, S_L is unbounded on the left and S_R is unbounded on the right. This is because (i) B_y is affine linear in y, and (ii) $B_y(L, S_L, S_R)$ is strictly decreasing in y, while $B_y(R, S_L, S_R)$ is strictly increasing in y. Thus, we have:

Corollary 1 *In an equilibrium the set of prosocial citizens who abstain is an interval (y_L, y_R) that contains $(x_L + x_R)/2$.*

The Corollary includes the possibility that one or both of the sets S_i are empty. Obviously, if voting costs are very high, only noise voters participate in equilibrium. We now derive an upper bound for c that ensures existence of non-trivial equilibria.

Because S_L and S_R take the forms $S_L = (-\infty, y_L]$ and $S_R = [y_R, \infty)$, to characterize the equilibrium it is sufficient to determine the values y_L and y_R, whom we refer to as the marginal voters of S_L and S_R, respectively. Because $\lambda = \beta = 0$ and $v_y(i) = -(x_i - y)^2$ we get:

$$B_y(L, (-\infty, y], S_R) = 2V\Phi(S_L)(x_R - x_L)\left(\frac{x_L + x_R}{2} - y\right);$$
$$B_y(R, [y, \infty), S_L) = 2V\Phi(S_R)(x_R - x_L)\left(y - \frac{x_L + x_R}{2}\right).$$

For the marginal voter y of S_i define $\gamma_i(y)$ to be y's gross (of cost) benefit of voting rather than abstaining:

$$\gamma_L(y) = 2V\Phi(y)(x_R - x_L)\left(\frac{x_L + x_R}{2} - y\right) - u(N);$$
$$\gamma_R(y) = 2V(1 - \Phi(y))(x_R - x_L)\left(y - \frac{x_L + x_R}{2}\right) - u(N).$$

Then, in equilibrium

$$\gamma_L(y_L) = \gamma_R(y_R) = c, \text{ and } S_L = (-\infty, y_L], \ S_R = [y_R, \infty). \tag{5}$$

Further, no minimal expansion yields that the slope of γ_i, $i = L, R$ must be negative at y_L and positive at y_R, i.e., $\gamma'_L(y_L) \le 0$ and $\gamma'_R(y_R) \ge 0$.

Assumption 4 (the log-concavity of ϕ) yields:

Lemma 1 *γ_L is single peaked in y on $(-\infty, (x_L + x_R)/2]$ and γ_R is single peaked in y on $[(x_L + x_R)/2, \infty)$.*

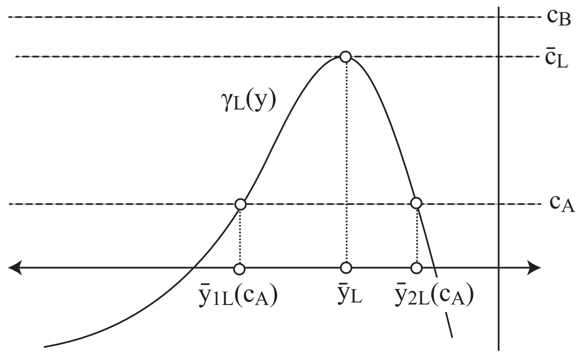

Figure 4: γ_L as a function of y for $u(N) = 0$ and the equilibrium solutions for different costs.

Let $\bar{c}_i$ be the maximum value of γ_i, $i = L, R$, i.e.,

$$\bar{c}_L = \max_{y<0.5(x_L+x_R)} \gamma_L(y), \text{ and } \bar{c}_R = \max_{y>0.5(x_L+x_R)} \gamma_R(y), \tag{6}$$

as illustrated in Figure 4. All prosocial i supporters abstain if $c > \bar{c}_i$. When $c = \bar{c}_i$, single-peakedness of γ_i implies that the equilibrium condition $\gamma_i(y) = c$ has a unique solution. For $c < \bar{c}_i$, Lemma 1 implies that the equilibrium equations (5) have two solutions. The derivative condition of no minimal expansion selects the root $\bar{y}_{2L}$ in Figure 4, which is the closest root to the mid-point between the two candidates. All citizens to the left of $\bar{y}_{2L}$ participate, while all other potential L supporters abstain. In contrast, no minimal expansion is violated at $\bar{y}_{1L}$. To see this, suppose that all citizens in $(\bar{y}_{1L}, \bar{y}_{1L} + dy]$ switch to vote for L. Then the net benefit of voter $\tilde{y} = \bar{y}_{1L} + dy$, given by the difference between γ_L and c_A in Figure 4, is strictly positive. Voters to the left of $\tilde{y}$ gain even more, because they care more about who wins, i.e., all deviating citizens are made strictly better off. Thus, a Nash equilibrium with $S_L = (-\infty, \bar{y}_{1L}]$ is not robust. Figure 4 also illustrates that if $c \le \bar{c}_i$ then turnout is bounded away from zero: participation for candidate L is at least $\Phi(\bar{y}_L)$.

It is immediate that y_L is strictly decreasing in c, while the opposite holds for y_R (see Figure 4). The set of types who abstain is given by $[y_L, y_R]$, which increases in c. In summary, we have established the following:

Proposition 1 *Non-trivial equilibria with strictly positive turnout of prosocial citizens for candidate i exist if and only if $c \le \bar{c}_i$. Any abstention is by moderates. If $c < \bar{c}_i$ is increased, then participation in the election strictly decreases.*

This result is in line with the empirical regularity documented in Section 2 that turnout decreases with the cost of voting.[19] It is immediate that an analogous result holds for the stakes, V, of the election—increases in V increase turnout.

[19]Section 7 extends the analysis to allow for continuous heterogeneity in voting costs. This extension does not just show that moderates abstain more than extremists, but that abstention is also single peaked as documented in Figure 3.

Continuity of the equilibrium in β implies that these results extend to β sufficiently small. Theoretically, there would be a non-monotonicity for very extreme types, e.g., very extreme liberals would abstain at a somewhat higher rate than extreme liberals, but for β sufficiently small these very extreme types are too many standard deviations away from the median to be observable in the data. Thus, when prosocial motives are sufficiently strong, turnout choices are consistent with the pattern of abstention documented in Figure 3.

6.2.3 Bandwagon vs Underdog Effects

The empirical literature documents the presence of bandwagon effects in two ways—(i) shifting the distribution of voter ideologies (e.g., by redistricting) and (ii) shifting voter beliefs about the distribution (e.g., by providing polling information to voters). The evidence shows that shifts that favor one candidate increase turnout propensities for that candidate and reduce turnout propensities for the rival candidate.

We now present our model analogues of these empirical findings for shifts that favor candidate R. To do this, we first shift the actual distribution of prosocial citizens from Φ to $\tilde{\Phi}$, such that $\tilde{\Phi}$ first-order stochastically dominates Φ, i.e., $\tilde{\Phi}(y) < \Phi(y)$, for all y in the interior of the supports. A particular example that we later use is a shift of the mean, i.e., $\tilde{\Phi}(y) = \Phi(y - m)$ for some $m > 0$. We then solely shift citizens' beliefs, fixing the true distribution of Φ. For example, this belief shift could be the result of a positive poll for candidate R. Let $\tilde{\Phi}^b$ be the shifted belief. This belief replaces Φ in γ_i, $i = L, R$, which we denote by $\tilde{\gamma}_i$. and thus determines $\tilde{S}_L$ and $\tilde{S}_R$. The actual numbers of prosocial voters are given by $\Phi(\tilde{S}_i)$, $i = L, R$.

Result 1 *Suppose that c is not so large that equilibria with positive participation for both candidates exist. Then:*

- ***Shifting the distribution of ideologies:*** *A first-order stochastic dominance shift of the ideology distribution from Φ to $\tilde{\Phi}$ reduces turnout for L and increases turnout for R: $\Phi(S_L) > \tilde{\Phi}(\tilde{S}_L)$ and $\Phi(S_R) < \tilde{\Phi}(\tilde{S}_R)$.*
- ***Shifting beliefs about distribution of ideologies:*** *Fixing Φ, a first-order stochastic dominance shift in beliefs about that distribution to $\tilde{\Phi}^b$ reduces turnout for L and increases turnout for R: $\Phi(\tilde{S}_L) < \Phi(S_L)$ and $\Phi(\tilde{S}_R) > \Phi(S_R)$.*

In both cases, the marginal prosocial voter for party i shifts to the left: $\tilde{y}_L < y_L < \frac{x_L + x_R}{2} < \tilde{y}_R < y_R$.

A shift of prosocial citizen ideologies to the right means that, all else equal, fewer such prosocial citizens vote for L and more vote for R. Further, the total effect goes beyond this direct effect, as the direct effect raises the prosocial benefit from voting for R-supporters and reduces it for L-supporters. This means that prosocial citizens must adjust their voting behavior to get back to equilibrium. Specifically, denoting the equilibrium

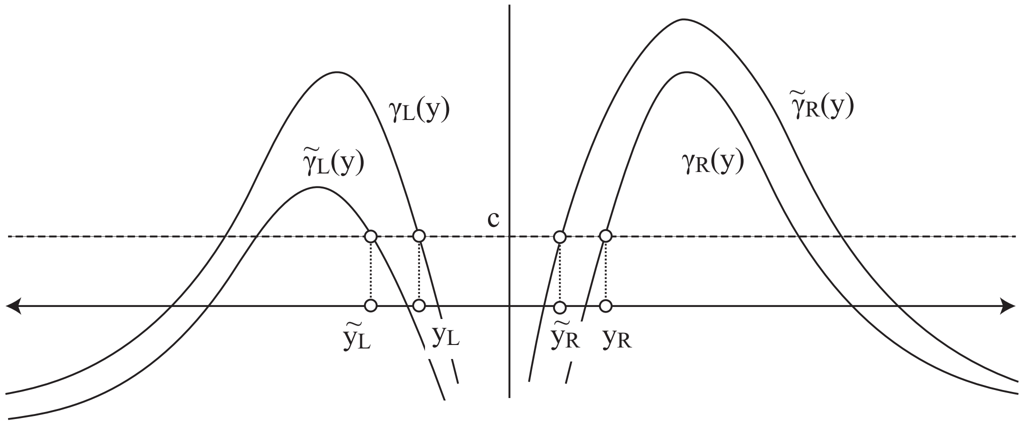

Figure 5: Effects of a first-order stochastic dominance shift of Φ to the right, shifting γ_i to $\tilde{\gamma}_i$, $i = L, R$.

cutoffs by y_L and y_R, no minimal expansion yields $\gamma'_L(y_L) < 0 < \gamma'_R(y_R)$. Thus, it must be that both y_L and y_R decrease, i.e., turnout for L falls even further, while turnout for R rises.

Inspection of Figure 5 makes this result transparent. The rightward shift of the distribution induces a uniformly downward shift in γ_L as $\Phi(y)$ decreases. Inspection yields that the new marginal voter $\tilde{y}_L$ for candidate L, given by the intersection of γ_L and c closest to $\frac{x_L+x_R}{2}$, must move to the left. Similarly, the rightward shift induces an upward shift in γ_R, because $1 - \Phi(y)$ increases and hence the intersection of γ_R and c closest to $\frac{x_L+x_R}{2}$ also moves to the left.

The second result reveals the effect of changing prosocial citizens' beliefs about candidate winning probabilities without affecting the preferences of voters. Now, the direct effect of shifting preferences is absent, but the indirect effect again induces a downward shift in γ_L and an upward shift in γ_R, and hence leftward shifts in their intersections with c closest to $\frac{x_L+x_R}{2}$.

Together, these findings reconcile the empirical evidence of bandwagon effects documented in Section 2.4.

6.2.4 Closeness of the Race

A prosocial citizen's strategic decision of whether or not to turn out hinges on how close the citizen believes the election will be. In our model, an election is a toss-up if the number of potential L and R supporters is about equal. When candidates are located symmetrically about zero, which is without loss of generality, the perfect toss-up case arises when the distribution Φ is symmetric around zero. From Result 1, we see that rightward shifts in the distribution of prosocial citizen ideologies, or rightward shifts in their beliefs about that distribution have opposite impacts on turnouts for candidates L and R. As a result, analytically characterizing impacts on overall turnout is more challenging.

This leads us to impose more structure. Below we consider rightward shifts in the distribution of prosocial ideologies of the form $\Phi_m(y) = \Phi(y - m)$: As m increases, the election becomes less competitive and safer for candidate R. We also consider the effect of fixing the distribution of ideologies while shifting the beliefs of prosocial citizens to the right so that beliefs are given by $\Phi_\eta(y) = \Phi(y - \eta)$, where $\eta > 0$ is the shift in beliefs in favor of candidate R, and the true distribution remains $\Phi(y)$.

We now identify simple conditions under which overall turnout is maximized by elections that are expected to be close.

Proposition 2 *Suppose that $x_R = -x_L$. Suppose that Φ is a uniform, normal or logistic distribution with a mean of zero. Then total turnout is locally maximized by a toss-up election, $m = 0$. Similarly, total turnout is locally maximized by beliefs that the election is a toss up, $\eta = 0$.*

To gain intuition for why an increase in m reduces turnout for L by more than it increases turnout for R, consider m positive but small. This implies that the distance between the marginal R voter, $y_R(m)$, and candidate L, located at x_L, is less than the distance between the marginal L voter, $y_L(m)$, and candidate R, located at x_R. This is because (i) for $m = 0$, the two distances are the same, i.e., $y_R(0) - x_L = x_R - y_L(0)$, and (ii) as shown above, both $y_R(m)$ and $y_L(m)$ decrease in m. Hence, with quadratic loss preferences, the relative proximity of $y_R(m)$ to x_L vis-à-vis the proximity of $y_L(m)$ to x_R means that the marginal R voter cares less about the victory of her preferred candidate than the marginal L voter. Increasing m thus has a larger direct effect on the net benefit of voting for the marginal L voter than the marginal R voter. The ideology of the marginal L voter must therefore adjust more to the increase in m than the ideology of the marginal R voter, i.e., $\frac{\partial y_L(m)}{\partial m} < \frac{\partial y_R(m)}{\partial m} < 0$. It follows that the width $y_R(m) - y_L(m)$ of the abstention region $[y_L(m), y_R(m)]$ increases. If Φ is uniform then it immediately follows that more voters abstain as m is moved away from zero. More generally, the shift of a strictly single peaked distribution reduces the mass in the abstention region, which in principle could offset the increased size of the abstention region. The proof provides a sufficient condition for this not to happen. The condition holds for distributions that do not have excessively fat tails, such as normal or logistic distributions. The sufficient condition fails, for example, for a t-distribution with one degree of freedom. Nonetheless, even for the t-distribution, we find numerically that turnout is still maximized at $m = 0$, i.e., at an ex-ante tossup election. For shifts only in voter beliefs, the result holds under weaker conditions as only the indirect effect discussed above matters and cannot possibly be offset by the shifts in the mass associated with a strictly single-peaked distribution.

These results are in line with the empirical regularity detailed in Section 2 that overall turnout is higher in closer elections.[20] Our model also provides a simple explanation for why the empirical magnitude of the impact of closeness on turnout is small: beliefs that the election will be closer move turnout for L and R in opposite directions. As a consequence, the effect on turnout is of second order.

[20] While Proposition 2 proves that participation decreases locally, we show numerically in Section 6.3 that the results holds globally and that we can quantitatively match the reduction in turnout.

6.2.5 Comparison With Other Voting Models

We now compare the performances of our model, in terms of fitting the empirical regularities discussed in Section 2, to existing voting models. We consider four alternative models:

Expressive Voting Model: Voters derive utility solely from the act of voting for a candidate with policy x. For example, with quadratic preferences, a citizen with ideology y votes instead of abstaining if $-(x - y)^2 - c \geq u(N)$. This corresponds to our model with purely selfish preferences, i.e., $\beta = 1$. Note that the voting patterns of this model are equivalent to those of models of abstention by indifference and alienation (e.g., Adams and Merrill (2003)). We pool them under the same header.

Abstention by Indifference: Citizens abstain from voting if and only if they are close to indifferent between the candidates, as in Erikson and Romero (1990). That is, given candidate positions x_L and x_R, citizen y abstains if and only if the utility gain from having their preferred candidate win is less than some constant $k > 0$, i.e., y abstains if $|-(x_L - y)^2 + (x_R - y)^2| \leq k$. Otherwise, y votes for their preferred candidate.

Pivotal Voting Models: A citizen votes if the expected utility from changing the electoral outcome due to their own participation (i.e., when pivotal), exceeds their voting cost (Palfrey and Rosenthal, 1983, 1985; Myerson, 1998; Castanheira, 2003; Myatt, 2015; Kawai et al., 2021; Miller, 2025). Maintaining quadratic preferences, the payoff of voting for, say, candidate L instead of abstaining for citizen θ equals Pr(pivotal for L)$\left((x_R - \theta)^2 - (x_L - \theta)^2\right) - c$. Kawai et al. (2021) and Miller (2025) allow the pivot probability to be larger (calling it "perceived voting efficacy"), allowing them to overcome the paradox of voting. But, all other comparative statics remain the same.

Group-Based Voting Models: In the classical ethical voting model (Feddersen and Sandroni, 2006; Coate and Conlin, 2004), a finite number of voters is divided ex-ante into being supporters of candidates L and R. Citizens who support candidate i have identical policy preferences, and only differ in their voting costs. Each group designs voting rules for their members that maximize the sum of expected utilities of group members, taking the actions of the other group as given. For some of our comparative static analyses, we must extend their model to a spatial setting, where supporters of a candidate differ in their policy preferences. Mobilization models, which assume that individuals' turnout decisions depend on the mobilization efforts exerted by leaders (Uhlaner, 1989; Morton, 1991; Shachar and Nalebuff, 1999), and social pressure models, which assume that individuals' decisions to vote are influenced by pressure from their peers (e.g., they want to appear prosocial or avoid being punished by their peers (Ali and Lin, 2013; Levine and Mattozzi, 2020)) produce similar comparative statics (for our purposes), and hence we pool them under the same header.

In an online appendix we detail the ways in which these alternative models can or cannot match the empirical regularities highlighted in Section 2. Table 2 summarizes these results and shows that our model outperforms these alternative models at matching those empirical regularities. Concretely, our model can reconcile all of them, whereas the alternative models fail to match at least 2 out of 5 of them.

	High Turnout in Safe and toss-up Elections	Raising voting costs reduces turnout	Higher turnout in closer elections	Bandwagon effect	Abstention is single-peaked in ideology
This model	✓	✓	✓	✓	✓
Expressive Voting	✓	✓	not always	×	×
Abstention by Indifference	✓	✓	×	×	✓
Pivotal Voting	×	✓	✓	×	✓
Group-Based Voting	×	✓	✓	×	×

Table 2: Comparison of model performances.

6.2.6 Additional Implications of Prosocial Voter Behavior

We next derive additional testable implications of our model, for which there is currently limited causal evidence in the empirical literature. Some of these predictions are shared by other models of voting.

We first consider how greater elite polarization in the form of increased separation between candidate platforms affects turnout:

Result 2 *Suppose that $c < \bar{c}_i$, $i = L, R$. Let $x_L < 0 < x_R$. Then for a fixed $x_L + x_R$, increasing elite polarization $x_R - x_L$ raises turnout.*

For $c < \bar{c}_R$, shifting x_R to the right and x_L to the left by the same amount increases $\gamma_R(y)$. Thus, y_R must shift to the left to retrieve $\gamma(y_R) = c$. This is because the marginal prosocial voter for R has more at stake when candidate platforms are further apart, and hence gains more from the costly efforts exerted by other voters for R. The analogous logic holds for candidate L. Thus, turnout rises.

The empirical literature provides suggestive evidence about the impact of elite polarization on turnout. More recent work (e.g., Béjar et al. (2020); Simas and Ozer (2021); Lehmann and Lindlacher (2026)) find positive estimated effects of polarization on turnout, while some earlier work (e.g., Rogowski (2014); Tausanovitch and Warshaw (2018)) find null or negative effects. Interpreting elite polarization as the stakes of the election for citizens, this result also speaks to the general pattern that turnout increases with the importance of the office for which candidates are running.[21]

[21] For instance, in the US, turnout is systematically higher in Presidential than in Senatorial or House elections (this is true both comparing across races for any given on-cycle election due to voter roll-off, and comparing across races in different electoral cycles, with turnout dropping substantially in mid-term elections).

We next consider the impact of increases in mass polarization on turnout:

Result 3 *Suppose $x_L = -x_R$ and Φ is symmetric. Let mass polarization be increased by replacing Φ with a symmetric distribution $\tilde{\Phi}$ that has more mass in the tails, i.e., $\tilde{\Phi}(y) > \Phi(y)$ for all $y < 0$. Then, turnout for each candidate strictly increases: $\tilde{\Phi}(\tilde{S}_i) > \Phi(S_i)$ for $i = L, R$.*

Shifting probability mass towards the tails implies that, ceteris paribus, γ_L shifts upwards, which, in turn, increases the cutoff y_L. The logic for candidate R is analogous.

We next consider the impact of one candidate becoming more extreme:

Result 4 *Suppose that $x_L < 0 < x_R$ and that x_R becomes more extreme while x_L does not change. Then turnout for candidate L rises, while turnout for candidate R can either rise or fall depending on parameters. That is, a more extreme candidate platform may or may not increase turnout for that candidate, but it always increases turnout for the rival.*

To see the result, we need to understand the impact of increasing x_R on the cutoffs y_L and y_R. As γ_L is increasing in x_R, y_L must increase. Intuitively, when x_R is more extreme, candidate R's victory is a worse outcome for L supporters, and this boosts their turnout due to the greater gains from the costly efforts exerted by other voters for L. In contrast, as Figure 6 illustrates, the impact on y_R is ambiguous. It is immediate that $\partial\gamma_R/\partial x_R$ is positive if $y > x_R$ (see the left panel) and negative if $y \leq x_R$ (see the right panel). Intuitively, candidate R's victory is a worse outcome for the marginal prosocial citizen y_R voting for R if a shift of x_R to the right moves x_R away from y_R, but it is a better outcome if it moves x_R closer to y_R.[22]

Result 4 unambiguously implies that a more extreme platform choice by one candidate always increases turnout for the other candidate, but the impact of the more radical platform choice on the candidate's own turnout can go either way. Lehmann and Lindlacher (2026) use the runoff structure of elections in Germany to provide consistent evidence: if an AfD candidate makes it to the runoff election then (1) there are large increases in turnout; (2) these turnout increases accrue solely to a rival who is a centrist; but (3) are split between the candidates when the rival is a left-winger. Further consistent evidence comes from Cutler et al. (2025), who use CES/CCES survey data about voters' perceptions of candidate ideology and find that more ideologically extreme candidates increase the turnout of both supporters and opponents.

[22] A naïve intuition might suggest that x_R is always larger than y_R because a prosocial citizen with ideal point x_R should want to be among those voting for R. However, when candidate policies are not very differentiated—when x_L and x_R are close—prosocial citizen $y = x_R$ may not care enough about who wins to benefit much from the voting efforts of others. In contrast, with quadratic loss preferences, prosocial citizens who are sufficiently extreme care more about who wins. If there are sufficiently many such extreme voters, i.e., if the standard deviation of Φ is sufficiently large, then these extreme prosocial citizens are incentivized to vote for R.

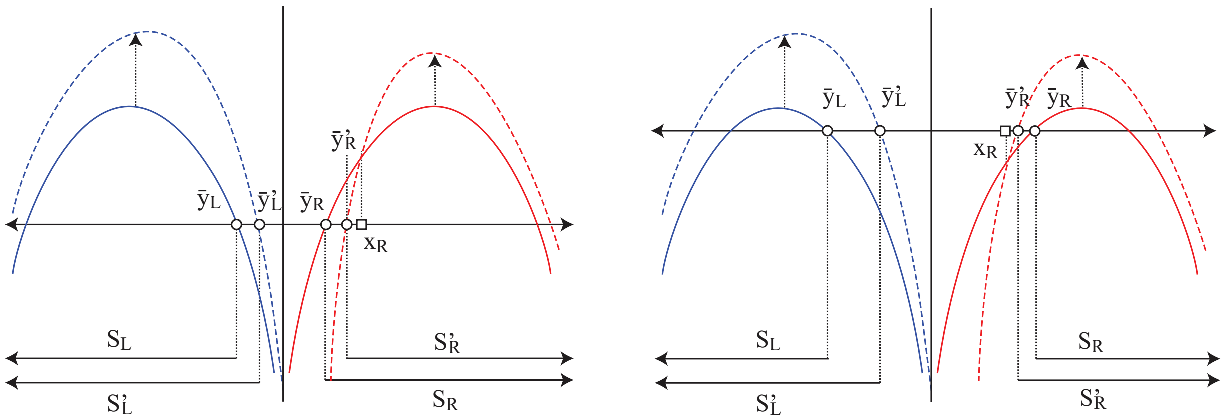

Figure 6: Increasing x_R increases turnout on the left, but may decrease (left panel) or increase (right panel) turnout on the right.

6.2.7 Impact of Adding Selfish Preferences

We next explore the impact adding selfish motives, so that the utility function becomes $u_y(i) = -\beta(x_i - y)^2 + (1-\beta)B_y(i;\cdot)$ with $0 < \beta < 1$. Theorem 1 proves that an equilibrium satisfying no minimal expansion exists. In particular, let

$$\gamma_{L,\beta}(y) = -\beta(x_L - y)^2 + 2(1-\beta)\left(\Phi(\bar{y}_L) - \Phi(\underline{y}_L)\right)(x_R - x_L)\left(\frac{x_L + x_R}{2} - y\right) - u(N). \tag{7}$$

Thus, if $S_L = [\underline{y}_L, \bar{y}_L]$ is the equilibrium participation set, then $\gamma_{L,\beta}(\underline{y}_L) = \gamma_{L,\beta}(\bar{y}_L) = c$. No minimal expansion implies that $\gamma'_{L,\beta}(\underline{y}_L) \geq 0$ and $\gamma'_{L,\beta}(\bar{y}_L) \leq 0$. Let $\beta' > \beta$. It is immediate that $\gamma_{L,\beta'}(y) < \gamma_{L,\beta}(y)$ for all citizen types y. Thus, $\underline{y}_L$ must weakly increase and $\bar{y}_L$ weakly decrease as β is increased. An analogous argument holds for S_R. Thus, we have:

Result 5 *Increases in prosocial motives,* $1-\beta$*, increase electoral participation.*

6.3 Quantitatively Matching Empirical Voting Patterns

In this section, we show that an appropriately-parametrized version of our model can quantitatively match established empirical voting patterns, even while assuming that all citizens share a common voting costs. In Section 7 we extend our model to allow for cost heterogeneity, which can only improve model fit.

Our quantitative exercise uses the gubernatorial election data discussed in Section 2.1. Recall that Figure 1 shows the average turnout in races rated as toss-up was 49.3% versus 41.2% for races rated as solid. Further, Figure 2 shows that on average, 15.5% of citizens still turn out to vote for candidates with no reasonable chance of winning—representing a substantial share of the lower total turnout in such elections. The win margins for races rated as toss-ups compared to races rated as solid for the Democrat or Republican candidate were 5.3% versus 29.5%, respectively.

We do not interpret the 8.1 percentage-point difference in average turnout between toss-up and solid races as the causal effect of electoral competitiveness. Some of this difference may reflect other factors outside the calculus of voting captured by our model, such as systematic differences in the underlying propensity to vote across states or electorates. We instead use the observed turnout difference as a demanding quantitative benchmark: our objective is to show that the model can generate variation in turnout even of this magnitude while simultaneously matching the other features of close and lopsided elections.

Thus, our model should deliver the following: (i) turnout is lower for races assessed as being safe, but (ii) it remains substantial even though (iii) there is negligible uncertainty about which candidate will win. So, too, (iv) the mean win margin in safe seats is 30%, and (v) the turnout for candidates who are almost certain to lose remains substantial. Finally, there is material uncertainty in both safe and toss-up races about the magnitudes of turnout and winning margins.

We now show that our model can *quantitatively* match the features of turnout in both close and lopsided elections. In addition to showing that we can match total turnout rates, we decompose total turnout to show that we can match separately the turnout rates for the underdog and frontrunner. To do this, we must specify the random turnout of noise voters on each side rather than just use the distribution over net noise voter support. In addition, as total turnout consists of both prosocial and noise voters, we must specify their relative shares of the pool of potential voters.

Recall that G measures the net turnout of noise voters for candidate R. We now model the actual distribution of turnout of noise voters for each candidate that determines net turnout. Let n_i, $i = L, R$ be the maximum number of possible noise voters for party i, and let ρ_i, $i = L, R$ be the random variable that describes the actual number who turn out. We assume that the ρ_is and Φ are normally distributed. Hence $\rho_R - \rho_L$ follows a normal distribution G. Then the turnout rate, T, in the polity and the win margin, W, are given by

$$T = \frac{\Phi(S_L) + \Phi(S_R) + \rho_L + \rho_R}{1 + n_L + n_R} \quad \text{and} \quad W = \left| \frac{\Phi(S_L) + \rho_L - \Phi(S_R) - \rho_R}{\Phi(S_L) + \rho_L + \Phi(S_R) + \rho_R} \right|. \tag{8}$$

Our comparative statics are on the impact of an ex-ante advantage for candidate R. In this analysis the average candidate position is at 0. Thus, citizens to the left of zero are potential L supporters, while those to the right are potential R supporters. The benchmark case where both candidates' ex-ante support is equally strong corresponds to Φ having mean zero, and ρ_L and ρ_R having the same mean. We then shift the mean of Φ and ρ_R to the right, and that of ρ_L to the left by the same margin, so that the base advantage for candidate R is the same both for prosocial and noise voters.

Figure 7 presents the sizes of $\Phi(S_L)$ and $\Phi(S_R)$ as a function of candidate R's ex-ante advantage. Note that turnout for L drops by more than the rise in turnout for R, implying that total turnout falls. Once R's advantage exceeds 17.5%, participation of prosocial L supporters drops to 0. This corresponds to the case in Figure 4 where $c > \bar{c}_L$. This discontinuity is an artifact of our simplifying assumption that all prosocial citizens have the same voting cost. In Section 7 we show that the discontinuity can disappear when we allow for a continuous distribution over voting costs.

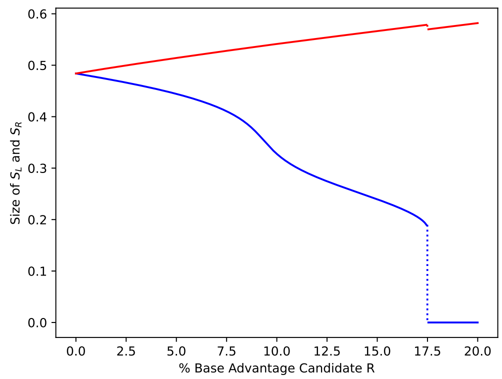

Figure 7: $\Phi(S_L)$ (lower curve) and $\Phi(S_R)$ (upper curve) as a function of candidate R's base advantage. Parameters: $c = 0.04$, $\lambda = 0.84$, $x_R = -x_L = 1$, $u_N = 0$. Both Φ and G are normal. The standard deviations of Φ and G are 1 and 0.1, respectively.

In Figures 8 and 9 we maintain the same parameters as in Figure 7, focusing on the range where the $\Phi(S_i)$ are continuous. We show that we can match quantitatively the features of lopsided elections documented above. We set $n_L = n_R = 0.8$, so that the maximum number of noise voters is 1.6 times the number of citizens with prosocial preferences and we consider independent normal distributions of noise voters with mean 0.2 and standard deviation $0.1/\sqrt{2}$ so the distribution G of $\rho = \rho_R - \rho_L$ has standard deviation 0.1.[23] The figures reveal substantial turnout obtains even when candidate L's winning probability is very close to zero. In the data, average turnout drops from 49.3% in toss-up elections to 41.2% in safe elections, i.e., an 8.1 percentage-point drop. In the figures, expected turnout drops from a high of 52.6% to 45.4% as the elections go from toss-up to safe, a 7.2 percentage-point decrease. In the data, the average win margins are 5.3% in toss-up elections and 29.5% in safe elections, while the associated margins are 5.9% and 33.4% for our model. Finally, in the most lopsided elections, the average turnout for the losing candidate 15.5% in the data and it is 15.2% in the model.

Numerical analyses reveal that λ primarily impacts the drop in turnout. In one extreme case, where $\lambda = 1$, all weight is placed on the electoral-impact component of prosocial preferences. In this case, support for the disadvantaged candidate drops to zero once the ex-ante disadvantage is even modestly large. This reflects that supporters of the disadvantaged candidate become discouraged as they have collectively only a very marginal impact on the electoral outcome. In contrast, if $\lambda = 0$, then turnout among prosocial citizens remains almost constant. Further, the shifts of the distributions of noise voters ρ_L and ρ_R do not matter, because they offset each other, leaving the distribution of total turnout of noise voters, $\rho_L + \rho_R$, unchanged. In fact, setting $\lambda = 0$ as we did in our theoretical analysis and keeping the other parameters unchanged, results in turnout dropping only from 52.58% to 52.56%, as we increase candidate R's ex-ante advantage from 0 to 17%.

[23]For ease of computation we do not truncate the normal distributions, as the probability of a realization outside the interval $[0, 0.8]$ is only about 0.002.

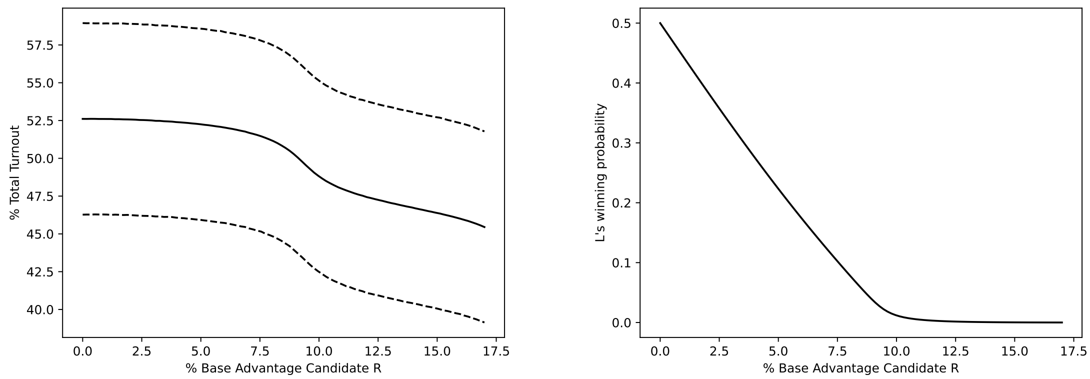

Figure 8: Turnout (left), and win probability (right), showing the average and the 90% confidence interval as a function of candidate *R*'s base advantage. Same parameters as in Figure 7.

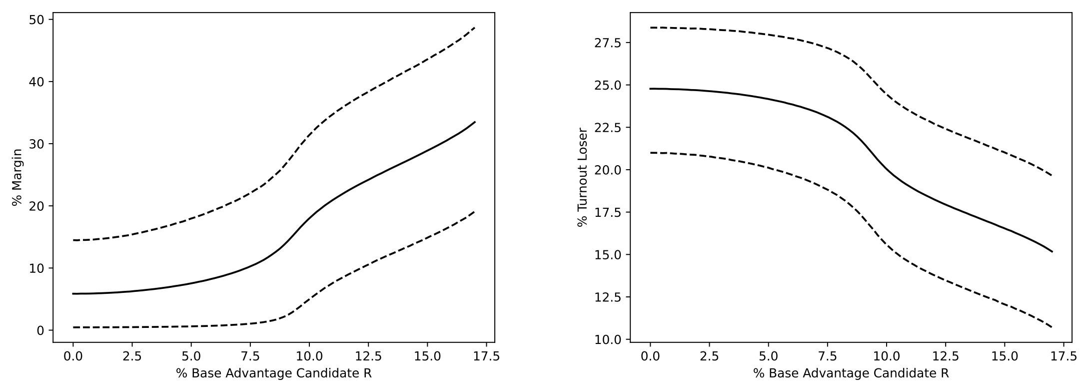

Figure 9: Winning margin (left) and turnout for losing candidate (right). Parameters as in Figure 7.

With our chosen parameter of $\lambda = 0.84$, we see a relatively large drop in turnout as the ex-ante advantage reaches about 5%. This is the point at which the impact of candidate L on the win probability becomes very small. This discouragement effect is dampened by the participation component of prosocial preferences, which is weighted by $1 - \lambda$: independently of their impact on the electoral outcome, greater participation by supporters of y's preferred candidate raises y's payoff from voting. This prevents an even more rapid decrease in prosocial voter support for an underdog who faces long odds of victory. Turnout is sensitive to the interplay between the two components, implying that we need roughly $5/6$ of the weight to be placed on the electoral-impact component to come close to matching the data.

Figure 9 shows how the endogenous responses of prosocial citizens quickly amplify a base electoral disadvantage into a major strategic disadvantage, which manifests itself in large realized turnout margin differences. Concretely, a 2% base disadvantage translates into an expected strategic disadvantage of 2.6% after accounting for the different turnout responses of the prosocial voters for the two candidates, while a 10%

base disadvantage results in a massive 20.4% strategic disadvantage that the underdog has almost no chance of overturning. This highlights that pollsters' forecasts of electoral outcomes should reflect not only the underlying balance of support, but also the endogenous turnout responses that amplify that initial advantage.

The 90% confidence intervals in the figures reveal that even the relatively small uncertainty about turnout of noise voters ($\sigma_G = 0.1$) has a significant impact on uncertainty about realized turnout and win margin. This uncertainty can be seen in the data on turnout and win margins for toss-up elections, where Cook Reports predict that the candidates have roughly equal chances of winning.

In sum, our simple calibration exercise can come close quantitatively to jointly matching the key features of how turnout varies with the ex-ante advantage of a favorite. Alternative parameterization of the model can readily deliver smaller dropoffs in turnout as elections go from tossup to safe, simply by reducing the value of λ. However, the decline of about 7% is close to the maximum of what our model can produce without generating a discontinuous drop in the support for the underdog. This limit reflects that we have a common voting cost. In the following section we show that much larger reductions in turnout can be generated with a continuous distribution over voting costs. A continuous distribution of voting costs is also necessary to generate the single peaked pattern of abstention as a function of ideology. Nonetheless, this analysis suggests that structural estimates of the model may deliver credible estimates of primitives. We leave that to future research.

7 Extension: Heterogeneous Voting Costs

To this point we have simplified the model and presentation by assuming that all citizens share the same voting cost. This specification turns out to be sufficient to explain the empirical facts that we documented. However, if one wants to take the model to the data more formally, one should allow for voting costs to differ among citizens. We now show that the model easily extends to incorporate cost heterogeneity, and that numerical solutions are readily obtained by solving a two-dimensional fixed-point problem.

With cost heterogeneity a prosocial type is characterized by (y, c), where y is the citizen's ideology and c is the voting cost. Let $\Phi(y, c)$ be the cdf describing the distribution of types. The definition of B_y is unchanged, but to compute it one now needs to integrate over both y and c to calculate the sizes of the sets of voters. By definition, a citizen's voting cost does not enter γ_i, $i = L, R$.

No minimal expansion, as detailed in Definition 3, must be modified to allow for an arbitrarily small set of citizens with different costs and ideologies to jointly deviate and vote for a candidate. In the one-dimensional case previously analyzed, the citizens who choose to jointly deviate have very similar ideologies to the marginal voter for a candidate. For example, if $S_L = (-\infty, y_L]$, then the citizens contemplating deviating to voting comprise the set $(y_L, y_L + dy)$. Alternatively, we could have defined a deviating set of citizens to be a collection of citizens whose utilities from unilaterally voting are very close to that of y_L. We now use this latter definition as it immediately generalizes to the two-dimensional setting in which citizens are distinguished by both ideology and cost.

When $\beta = 0$, two straight lines in the ideology-cost space separates voters from non-voters, one for candidate L, the other for candidate R. Let $\bar{x} = (x_L + x_R)/2$. Both lines of separation go through $(\bar{x}, -u_N)$ because citizens with that ideology are indifferent between the candidates and vote if and only if $-c > u_N$, i.e., all citizens $(\bar{x}, c)$ with costs $c < -u_N$ vote. The line that determines the set of R voters has positive slope, and the line for L voters has a negative slope. Let (x, c) be on one of the lines, and consider citizen $(x, c + dc)$, where $x \neq \bar{x}$. The citizens on the same indifference curve are given by a straight line through $(x, c+dc)$ and $(\bar{x}, -u_N)$. No minimal expansion means that these citizens and all those with stronger preferences for the candidate who previously did not vote for the candidate cannot all be made better off by voting. The γ_i functions in this extension are therefore defined by the slope of the line that separates candidate i voters from non-voters, thus reducing the search for an equilibrium to finding the solution of a two-dimensional fixed point problem as in the single dimensional setting when G is not uniform. Numerical solution is straightforward. The existence proof extends to the case of continuous costs by starting with a grid over both ideology and costs. The limiting argument over c works because of the monotonicity of the set of abstainers in c.

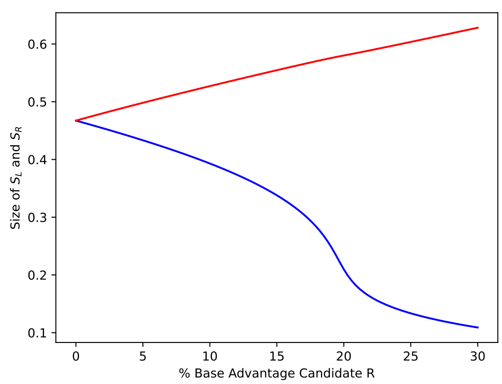

Figure 10: Sizes of S_L and S_R as a function of candidate R's base advantage. Costs are uniformly distributed on $[0, 0.5]$. $u_N = -0.1$, and $\sigma_G = 0.2$. All other parameters are as in Section 6.3.

Figure 10 shows that, unlike in the single cost setting, voter participation can be a continuous function of a candidate's base advantage. Of course, if the standard deviation of G is too small or there is too little cost uncertainty, then support for the disadvantaged candidate can collapse discontinuously as in the single-cost setting. Note that, consistent with our earlier theoretical analysis, increasing candidate R's advantage has a discouragement effect on L supporters that exceeds the encouragement effect on R supporters, i.e., total turnout is locally maximized in toss-up elections.

In the single-cost setting, raising voting costs reduces turnout. This is clearly still true with heterogeneous costs, but more nuanced phenomena also emerge: Raising costs for some ideology types not only reduces their participation, but it also has a secondary effect of reducing participation of all other y types who support the same candidate. To understand the real-world consequence, suppose a state makes voting more difficult in an urban district predominated by Democratic voters. The direct effect is to reduce turnout in that urban district, but this, in turn, also discourages potential Democratic voters in otherwise unaffected districts due to the reduced winning probability and overall turnout that enter the prosocial component of preferences.

Allowing for heterogeneity in voting costs immediately generates a single-peaked relationship between abstention and ideology for prosocial citizens when y and c are independently distributed: The equilibrium condition (5) directly implies that the interval of centrists who abstain is strictly increasing in c, unless costs are so prohibitively large that all ideology types with such voting costs abstain. Integrating over c yields that abstention is single peaked as a function of ideology.

8 Conclusion

Voting in a large election has a public good aspect, as the decisions of some citizens to participate provide a positive externality to other citizens who prefer the same candidate. A vast empirical literature has established that people behave far more pro-socially than theoretical models based on purely selfish behavior would predict. We propose a model of voting in which citizens have prosocial motives, making them averse to free riding on the costly voting efforts of co-supporters. Our model integrates three distinct aspects of prosocial considerations. First, citizens are more willing to vote when more supporters of their preferred candidate do so as well, independently of the impact of their participation on the electoral outcome. Second, citizens are more willing to vote when the collective voting efforts of these supporters have a larger impact on their preferred candidate's winning probability. Third, the prosocial payoff from voting is higher for citizens who care more about which candidate wins the election.

We document several core empirical features of turnout and show that our model can reconcile all of them, whereas existing models are inconsistent with several of these features. We also show that each of the three components of prosocial preferences is needed to quantitatively match the data. The first component is necessary to sustain substantial turnout for a candidate who has almost no chance of winning. The second is essential for generating the decline in total turnout as elections become less competitive. The third—that the prosocial payoff from voting is larger for citizens who care more about which candidate wins—is essential for explaining how turnout varies with ideology and for preventing turnout from approaching 100

Our model is tractable and can be integrated into other political economy models, for example, to investigate political competition among candidates. Moreover, model solutions are Nash equilibria, and they exist under very general conditions.

References

Abramson, P. R., J. H. Aldrich, A. Diskin, A. M. Houck, R. Levine, T. J. Scotto, and D. B. Sparks (2018). The effect of national and constituency expectations on tactical voting in the British general election of 2010. In *The many faces of strategic voting*, pp. 28. University of Michigan Press Ann Arbor, MI.

Adams, J. and S. Merrill (2003). Voter turnout and candidate strategies in American elections. *Journal of Politics 65*, 161–189.

Agranov, M., J. K. Goeree, J. Romero, and L. Yariv (2018). What makes voters turn out: The effects of polls and beliefs. *Journal of the European Economic Association 16*(3), 825–856.

Ainsworth, R., E. G. Munoz, A. M. Gomez, et al. (2024). District competitiveness increases voter turnout: evidence from repeated redistricting in North Carolina. *Quarterly Journal of Political Science 19*(4), 387–432.

Ali, S. N. and C. Lin (2013, May). Why people vote: Ethical motives and social incentives. *American Economic Journal: Microeconomics 5*(2), 73–98.

Andreoni, J. and D. Bernheim (2009). Social image and the 50-50 norm: A theoretical and experimental analysis of audience effects. *Econometrica 77*, 1607–1636.

Andreoni, J. and J. Miller (2002). Giving according to GARP: An experimental test of the consistency of preferences for altruism. *Econometrica 70*, 737–753.

Bagwe, G., J. Margitic, and A. Stashko (2022). Polling place location and the costs of voting. Technical report, Working Paper.

Battigalli, P. and M. Dufwenberg (2007). Guilt in games. *American Economic Review 97*, 170–176.

Battigalli, P. and M. Dufwenberg (2009). Dynamic psychological games. *Journal of Economic Theory 144*, 1–35.

Béjar, S., J. A. Moraes, and S. López-Cariboni (2020). Elite polarization and voting turnout in latin america, 1993–2010. *Journal of Elections, Public Opinion and Parties 30*(1), 1–21.

Belli, R., M. Traugott, M. Young, and K. McGonagle (1999). Reducing vote overreporting in surveys: Social desirability, memory failure, and source monitoring. *Public Opinion Quarterly 63*, 90–108.

Blais, A. (2000). *To vote or not to vote: The merits and limits of rational choice theory*. University of Pittsburgh Press.

Bolton, G. and A. Ockenfels (2000). ERC: a theory of equity, reciprocity, and competition. *American Economic Review 90*, 166–193.

Bond, R. M., C. J. Fariss, J. J. Jones, A. D. Kramer, C. Marlow, J. E. Settle, and J. H. Fowler (2012). A 61-million-person experiment in social influence and political mobilization. *Nature 489*(7415), 295–298.

Bouton, L., M. Castanheira, and L.-S. Aniol (2017). Multicandidate elections: Aggregate uncertainty in the laboratory. *Games and Economic Behavior 101*, 132–150.

Braconnier, C., J.-Y. Dormagen, and V. Pons (2017). Voter registration costs and disenfranchisement: experimental evidence from france. *American Political Science Review 111*(3), 584–604.

Brady, H. E. and J. E. McNulty (2011). Turning out to vote: The costs of finding and getting to the polling place. *American Political Science Review 105*(1), 115–134.

Bursztyn, L., D. Cantoni, P. Funk, F. Schönenberger, and N. Yuchtman (2024). Identifying the effect of election closeness on voter turnout: evidence from swiss referenda. *Journal of the European Economic Association 22*(2), 876–914.

Bénabou, R. and J. Tirole (2006). Incentives and prosocial behavior. *American Economic Review 96*, 1652–1678.

Cantoni, E. (2020). A precinct too far: Turnout and voting costs. *American Economic Journal: Applied Economics 12*(1), 61–85.

Castanheira, M. (2003). Victory margins and the paradox of voting. *European Journal of Political Economy 19*(4), 817–841.

Charness, G. and M. Rabin (2002). Understanding social preferences with simple tests. *Quarterly Journal of Economics 117*, 817–869.

Chaudhuri, A. (2011). Sustaining cooperation in laboratory public goods experiments: a selective survey of the literature. *Experimental Economics 14*, 47–83.

Clinton, J. D., N. Eubank, A. Fresh, and M. E. Shepherd (2021). Polling place changes and political participation: evidence from North Carolina presidential elections, 2008–2016. *Political Science Research and Methods 9*(4), 800–817.

Coate, S. and M. Conlin (2004). A group rule-utilitarian approach to voter turnout: Theory and evidence. *American Economic Review 94*(5), 1476–1504.

Coma, F. M. i. and D. L. Van De Maele (2023). The global dataset on turnout (gd-turnout). *Electoral Studies 86*, 102681.

Concha-Arriagada, C. and J. Naddeo (2022). It's always sunny in politics. Working Paper.

Cubitt, R. P., M. Drouvelis, S. Gächter, and R. Kabalin (2011). Moral judgments in social dilemmas: How bad is free riding? *Journal of Public Economics 95*(3-4), 253–264.

Cutler, A. L., H. J. Hassell, and K. Reuning (2025). The effect of candidate ideological extremism on vote choice and turnout. Working Paper.

Daoust, J.-F. (2018). Support for minority government and strategic voting. In J. H. Aldrich, A. Blais, and L. B. Stephenson (Eds.), *The many faces of strategic voting*, Chapter 4, pp. 75–88. University of Michigan Press Ann Arbor, MI.

DellaVigna, S., J. A. List, U. Malmendier, and G. Rao (2017). Voting to tell others. *Review of Economic Studies 84*(1), 143–181.

Doherty, D., C. Dowling, A. Gerber, and G. Huber (2017). Are voting norms conditional? How electoral context and peer behavior shape the social returns to voting. *Journal of Politics 79*, 1095–1100.

Duffy, J. and M. Tavits (2008). Beliefs and voting decisions: A test of the pivotal voter model. *American Journal of Political Science 52*(3), 603–618.

Enos, R. D. and A. Fowler (2014). Pivotality and turnout: Evidence from a field experiment in the aftermath of a tied election. *Political Science Research and Methods 2*(2), 309–319.

Erikson, R. S. and D. W. Romero (1990). Candidate equilibrium and the behavioral model of the vote. *American Political Science Review 84*(4), 1103–1126.

Esponda, I. and E. Vespa (2014). Hypothetical thinking and information extraction in the laboratory. *American Economic Journal: Microeconomics 6*(4), 180–202.

Falk, A. and U. Fischbacher (2006). A theory of reciprocity. *Games and Economic Behavior 54*, 293–315.

Feddersen, T. and A. Sandroni (2006). A theory of participation in elections. *American Economic Review 96*(4), 1271–1282.

Fehr, E. and S. Gächter (2002). Altruistic punishment in humans. *Nature 415*(6868), 137–140.

Fehr, E. and S. Gächter (2000). Cooperation and punishment in public goods experiments. *American Economic Review 90*, 980–994.

Fehr, E. and K. M. Schmidt (1999). A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics 114*(3), 817–868.

Fehr, E. and I. Schurtenberger (2018). Normative foundations of human cooperation. *Nature Human Behaviour 2*, 458–468.

Fischbacher, U., S. Gächter, and E. Fehr (2001). Are people conditionally cooperative? Evidence from a public goods experiment. *Economics Letters 71*(3), 397–404.

Forsythe, R., T. Rietz, R. Myerson, and R. Weber (1996). An experimental study of voting rules and polls in three-candidate elections. *International Journal of Game Theory 25*, 355–383.

Fowler, J. (2006). Altruism and turnout. *Journal of Politics 68*, 674–683.

Fraga, B. L., D. J. Moskowitz, and B. Schneer (2022). Partisan alignment increases voter turnout: Evidence from redistricting. *Political Behavior 44*(4), 1883–1910.

Fujiwara, T. (2011). A regression discontinuity test of strategic voting and duverger's law. *Quarterly Journal of Political Science 6*, 197–233.

Funk, P. (2010). Social incentives and voter turnout: Evidence from the Swiss mail ballot system. *Journal of the European Economic Association 8*(5), 1077–1103.

Gachter, S. and B. Herrmann (2009). Reciprocity, culture and human cooperation: previous insights and a new cross-cultural experiment. *Philosophical Transactions of the Royal Society 364*, 791–806.

Gerber, A., M. Hoffman, J. Morgan, and C. Raymond (2020, July). One in a million: Field experiments on perceived closeness of the election and voter turnout. *American Economic Journal: Applied Economics 12*(3), 287–325.

Gerber, A. and T. Rogers (2009). Descriptive social norms and motivation to vote: Everybody's voting and so should you. *Journal of Politics 71*, 178–191.

Gerber, A. S., D. P. Green, and C. W. Larimer (2008). Social pressure and voter turnout: Evidence from a large-scale field experiment. *American Political Science Review 102*(1), 33–48.

Gerber, A. S., G. A. Huber, D. Doherty, and C. M. Dowling (2016). Why people vote: Estimating the social returns to voting. *British Journal of Political Science 46*(2), 241–264.

Gerber, A. S., G. A. Huber, A. H. Fang, and A. Gooch (2017). The generalizability of social pressure effects on turnout across high-salience electoral contexts: Field experimental evidence from 1.96 million citizens in 17 states. *American Politics Research 45*(4), 533–559.

Gomez, B. T., T. G. Hansford, and G. A. Krause (2007). The republicans should pray for rain: Weather, turnout, and voting in US presidential elections. *Journal of Politics 69*(3), 649–663.

Großer, J. and A. Schram (2010). Public opinion polls, voter turnout, and welfare: An experimental study. *American Journal of Political Science 54*(3), 700–717.

Harris, A., H. Walker, A. White, J. Doleac, L. Eckhouse, and E. Foster-Moore (2025). Registering returning citizens to vote: A field experiment in North Carolina. Working Paper.

Jankowski, R. (2007). Altruism and the decision to vote. *Rationality and Society 19*, 5–34.

Kaplan, E. and H. Yuan (2020, January). Early voting laws, voter turnout, and partisan vote composition: Evidence from Ohio. *American Economic Journal: Applied Economics 12*(1), 32–60.

Kartal, M. (2015). Laboratory elections with endogenous turnout: Proportional representation versus majoritarian rule. *Experimental Economics 18*, 366–384.

Kawai, K., Y. Toyama, and Y. Watanabe (2021). Voter turnout and preference aggregation. *American Economic Journal: Microeconomics 13*(4), 548–586.

Kelley, H. H. and A. J. Stahelski (1970). Social interaction basis of cooperators' and competitors' beliefs about others. *Journal of Personality and Social Psychology 16*(1), 66.

Kimbrough, E. O. and A. Vostroknutov (2016). Norms make preferences social. *Journal of the European Economic Association 14*, 608–638.

Kiss, H. J., A. Rosa-Garcia, and V. Zhukova (2020). Conditional cooperation in group contests. *PLoS ONE 15*(12), e0244152.

Knack, S. (1992). Civic norms, social sanctions, and voter turnout. *Rationality and Society 4*, 133–156.

Kocher, M., T. Cherry, S. Kroll, R. Netzer, and M. Sutter (2008). Conditional cooperation on three continents. *Economics Letters 101*, 175–178.

Lehmann, K. and V. Lindlacher (2026). Far-right entry and voter mobilization in local runoff elections. Working Paper, ifo Institute of Economic Research.

Levine, D. (1998). Modeling altruism and spitefulness in experiments. *Review of Economic Dynamics 1*, 593–622.

Levine, D. and A. Mattozzi (2020). Voter turnout with peer punishment. *American Economic Review 110*(10), 3298–3314.

Malmendier, U., V. L. te Velde, and R. A. Weber (2014). Rethinking reciprocity. *Annual Review of Economics 6*(1), 849–874.

Miller, L. (2025). Motivated or mobilized? Competitiveness, campaign effort, and turnout in U.S. elections. Working Paper.

Morton, R. (1991). Groups in rational turnout models. *American Journal of Political Science 35*(3), 758–776.

Moskowitz, D. J. and B. Schneer (2019). Reevaluating competition and turnout in US House elections. *Quarterly Journal of Political Science 14*(2), 191–223.

Myatt, D. P. (2015). A theory of voter turnout. Working Paper.

Myerson, R. B. (1998). Population uncertainty and Poisson games. *International Journal of Game Theory 27*, 375–392.

Opp, K.-D. (2001). Why do people vote? the cognitive-illusion proposition and its test. *Kyklos 54*, 355–378.

Ortoleva, P. and E. Snowberg (2015). Overconfidence in political behavior. *American Economic Review 105*(2), 504–535.

Palfrey, T. R. and H. Rosenthal (1983). A strategic calculus of voting. *Public Choice 41*, 7–53.

Palfrey, T. R. and H. Rosenthal (1985). Voter participation and strategic uncertainty. *American Political Science Review 79*, 62–78.

Pons, V. and C. Tricaud (2018). Expressive voting and its cost: Evidence from runoffs with two or three candidates. *Econometrica 86*(5), 1621–1649.

Rabin, M. (1993). Incorporating fairness into game theory and economics. *American Economic Review 83*, 1281–1302.

Riker, W. H. and P. C. Ordeshook (1968). A theory of the calculus of voting. *American Political Science Review 62*(1), 25–42.

Rogowski, J. C. (2014). Electoral choice, ideological conflict, and political participation. *American Journal of Political Science 58*(2), 479–494.

Rotemberg, J. (2008). Minimally acceptable altruism and the ultimatum game. *Journal of Economic Behavior and Organization 66*, 457–476.

Shachar, R. and B. Nalebuff (1999). Follow the leader: Theory and evidence on political participation. *American Economic Review 89*(3), 525–547.

Silver, B., B. Anderson, and P. Abramson (1986). Who overreports voting? *American Political Science Review 80*, 613–624.

Simas, E. N. and A. L. Ozer (2021). Polarization, candidate positioning, and political participation in the US. *Electoral Studies 73*, 102370.

Spenkuch, J. L. (2018). Expressive vs. strategic voters: An empirical assessment. *Journal of Public Economics 165*, 73–81.

Tausanovitch, C. and C. Warshaw (2018). Does the ideological proximity between candidates and voters affect voting in US house elections? *Political Behavior 40*, 223–245.

Thompson, D. M., J. A. Wu, J. Yoder, and A. B. Hall (2020). Universal vote-by-mail has no impact on partisan turnout or vote share. *Proceedings of the National Academy of Sciences 117*(25), 14052–14056.

Tomkins, S., K. Yao, J. Gaebler, T. Konitzer, D. Rothschild, M. Meredith, and S. Goel (2023). Blocks as geographic discontinuities: The effect of polling-place assignment on voting. *Political Analysis 31*(2), 165–180.

Uhlaner, C. J. (1989). Rational turnout: The neglected role of groups. *American Journal of Political Science 33*(2), 390–422.

Yoder, J. (2018). How polling place changes reduce turnout: Evidence from administrative data in North Carolina. Working Paper.

Yoder, J., C. Handan-Nader, A. Myers, T. Nowacki, D. M. Thompson, J. A. Wu, C. Yorgason, and A. B. Hall (2021). How did absentee voting affect the 2020 US election? *Science Advances 7*(52), eabk1755.

Appendix: Proofs

Proof of Theorem 1. For every $n \in \mathbb{N}$ with $n \geq 1$, let A_i^n, $i = 1, \ldots, n$, be a partition of $\bar{\mathbb{R}} = [-\infty, \infty]$ into intervals of the same measure, i.e., $\Phi(A_i^n) = 1/n$ for all $i = 1, \ldots, n$.

Let Z^n be the collection of midpoints of the n^{th} partition. For $S \subset Z^n$ define $\psi_n(S) = \bigcup_{\{i|z \in S\}} A_i^n$ to be the set in $\bar{\mathbb{R}}$ that consists of all partition elements containing points in S. Define $\tilde{B}_y(i, S_L, S_R) = B_y(i, \psi_n(S_L), \psi_n(S_R))$. Thus, $\tilde{B}_y$ and hence utility $\tilde{u}_y = v_y(\cdot) + \tilde{B}_y(\cdot)$ is defined on Z^n. As in the continuum case, each $z \in Z^n$ is a prosocial citizen who decides whether or not to vote. As Z^n and the players' strategy sets are finite, there exists a mixed strategy Nash equilibrium of the game.

Let S_i^n be the set of all prosocial citizens in Z^n who vote for candidate $i \in \{L, R\}$ with strictly positive probability. First suppose that z votes for R and z' for candidate L, where $z < z'$. Then the assumption that the net benefit of voting for R instead of L is strictly increasing in y implies that z' should vote for R if z votes for R, a contradiction.

Next, suppose that there exist $z < z' < z''$ such that z and z'' vote with positive probability for L, but z' either does not vote, or votes with a probability strictly less than 1. There exists $\alpha \in (0, 1)$ such that $z' = \alpha z + (1 - \alpha) z''$. Let $\tilde{A} = \psi_n(\{z'\})$ be the element of the partition that contains z'. Note that $\Phi(\tilde{A}) > 0$. Define $T_i^n(\omega)$, $\omega \in \Omega$, to be the realized set of citizens in the discrete model excluding z' who vote for candidate i. Thus, $\tilde{A} \cap \psi_n(T_L^n(\omega)) = \emptyset$. Then

$$\begin{aligned}
E_\Omega\left[\tilde{B}_{z'}(L, T_L^n(\omega) \cup \{z'\}, T_R^n(\omega))\right] &= E_\Omega\left[B_{z'}(L, \psi_n(T_L^n(\omega)) \cup \tilde{A}, \psi_n(T_R^n(\omega)))\right] \\
&\geq E_\Omega\left[B_{z'}(L, \psi_n(T_L^n(\omega)), \psi_n(T_R^n(\omega)))\right] \\
&\geq \alpha E_\Omega\left[B_z(L, \psi_n(T_L^n(\omega)), \psi_n(T_R^n(\omega)))\right] + (1 - \alpha) E_\Omega\left[B_{z''}(L, \psi_n(T_L^n(\omega)), \psi_n(T_R^n(\omega)))\right] \\
&= \alpha E_\Omega\left[\tilde{B}_z(L, T_L^n(\omega), T_R^n(\omega))\right] + (1 - \alpha) E_\Omega\left[\tilde{B}_{z''}(L, T_L^n(\omega), T_R^n(\omega))\right],
\end{aligned}$$

where the first inequality follows from monotonicity and the second inequality from concavity of B_y in y. Further, by assumption, at least one of the inequalities is strict, because B_y is either strictly monotone or strictly concave. Thus, because v_y is concave in y, utility $\tilde{u}_y$ is strictly concave in y. Hence,

$$\begin{aligned}
E_\Omega\left[\tilde{u}_{z'}(L, T_L(\omega), T_R(\omega), x_L, x_R)\right] &> \alpha E_\Omega\left[\tilde{u}_z(L, T_L(\omega), T_R(\omega), x_L, x_R)\right] + (1 - \alpha) E_\Omega\left[\tilde{u}_{z''}(L, T_L(\omega), T_R(\omega), x_L, x_R)\right] \\
&\geq \alpha u(N) + (1 - \alpha) u(N) = u(N).
\end{aligned}$$

Thus, prosocial citizen type z' is strictly better off voting for L than abstaining. Thus, z' must vote with probability 1. This implies that each set S_L^n cannot have gaps, and that voting with strictly interior probability can only occur at the endpoints. The same holds for all sets S_R^n. Hence, these sets are intervals.

Let $(S_L^n)_{n \in \mathbb{N}}$, $(S_R^n)_{n \in \mathbb{N}}$ be the sequence of such intervals. Because $\bar{\mathbb{R}}$ is compact, it follows that there exist subsequences that converge to sets S_L and S_R. With a slight abuse of notation, denote these subsequences again by $(S_L^n)_{n \in \mathbb{N}}$ and $(S_R^n)_{n \in \mathbb{N}}$. We next prove S_L, S_R is a Nash equilibrium of the continuum model.

Suppose by way of contradiction that S_L and S_R do not comprise a Nash equilibrium. Let $y \in S_L$ and suppose that $u_y(L, S_L, S_R) < u(N)$. Because $\phi > 0$ on its interval support, the length of the intervals in each partition must go to zero as $n \to \infty$. Thus, there exists a sequence of points $z_n \in Z^n$ with $z_n \to y$. However, because $u_{z_n}(L, \psi_n(S^n_L), \psi_n(S^n_R)) \geq u(N)$ for all n, and the sets $\psi_n(S^n_j)$, $j = L, R$ are intervals converging to S_L and S_R, continuity (implied by Assumption 2) implies that $u_y(L, S_L, S_R) \geq u(N)$, a contradiction. Now let $y \notin S_L$ and suppose that $u_y(L, S_L, S_R) > u(N)$. Then $u_{z_n}(L, S^n_L, S^n_R) > u(N)$ for all sufficiently large n. But then $z_n \in S^n_L$ for all large n, and hence $y \in S_L$, a contradiction.

Because prosocial citizens strictly outside S_L are worse off voting for L and prosocial citizens in S_L are weakly better off, this implies that the prosocial citizens at the boundary points of S_L are indifferent between voting and abstaining. If they mix, we can purify their strategies and assume that they either vote with probability 1 or abstain with probability 1. The utilities of all other citizens do not change by Assumption 3.

Finally, we need to prove that no minimal expansion holds in the limit. Intuitively, this follows because the property of no minimal expansion extends to the continuum model the feature that, in a finite game, each individual has a non-negligible effect on winning probabilities. Suppose by way of contradiction that no minimal expansion is violated, without loss of generality, for S_L. As S_L is an interval, denote its lower bound by $\underline{y}_L$ and its upper bound by $\bar{y}_L$. Suppose, for example, that no minimal expansion is violated at the upper bound, i.e.,

$$\frac{\partial u_y(L, [\underline{y}_L, y], S_{-i})}{\partial y}\Big|_{y=\bar{y}_L} > 0.$$

By assumption, this derivative is continuous in a neighborhood of $\bar{y}_L$. Hence, there exists a $\delta > 0$ such that

$$\frac{\partial u_y(L, [y_1, y], S'_{-i})}{\partial y}\Big|_{y=y_2} > 0,$$

for all $y_1 \in [\underline{y}_L - \delta, \underline{y}_L + \delta]$, $y_2 \in [\bar{y}_L - \delta, \bar{y}_L + \delta]$ and all intervals S'_{-i} close to S_{-i}. Let $S^i_n = [\underline{y}^n_i, \bar{y}^n_i]$. For every S^n_i, let $z_n = \min\{z \in Z^n | z > \bar{y}^n_i\}$. Then

$$\frac{\partial u_y(L, [\underline{y}^n_i, y], S^n_{-i})}{\partial y}\Big|_{y=z_n} > 0,$$

for n sufficiently large. This means that prosocial citizen z_n is strictly better off if all prosocial citizens in $[\bar{y}^n_i, z_n]$ vote for candidate i instead of abstaining. Because $\psi_n(z_n) \supset [\bar{y}^n_i, z_n]$, prosocial citizen z_n is also strictly better off if all prosocial citizens in $\psi_n(z_n)$ vote for candidate i. However, this means that in the finite game where prosocial citizens are represented by points in Z^n, prosocial citizen z^n could obtain a strictly higher payoff by voting for candidate i rather than abstaining, a contradiction. Analogous arguments establish the result at $\underline{y}_L$. Thus, in the equilibrium, no minimal expansion is satisfied.

The second statement follows immediately from the concavity argument used above. That is, $z < z' < z''$ such that $z, z'' \in S_i$, but z' either does not vote, or votes with a probability strictly less than 1. Then by concavity and monotonicity of utility it follows that z' must strictly prefer to vote. ■

Proof of Lemma 1. Let $\bar{x} = (x_L + x_R)/2$. Differentiating γ_L with respect to y yields

$$2V\phi(y)(x_R - x_L))\,(\bar{x} - y) - 2V\Phi(y)(x_R - x_L). \tag{9}$$

The sign of this expression does not change if we divide it by $2V\Phi(y)(x_R - x_L)$. Doing so yields

$$\frac{\phi(y)}{\Phi(y)}\,(\bar{x} - y) - 1. \tag{10}$$

Log concavity implies that $\phi(y)/\Phi(y)$ is decreasing in y. Given that $(\bar{x} - y)$ is also decreasing in y, we have that (10) is decreasing in y. Further, this expression is positive for sufficiently negative y, and strictly negative at $y = \bar{x}$. Thus, (10) has exactly one root, which implies that (9) has exactly one root. Further, (9) is negative at $y = \bar{x}$. Hence, the unique root is strictly less than $\bar{x}$ and is the unique local maximizer of γ_L. The argument for γ_R is analogous. ■

Proof of Proposition 1. See text. ■

Proof of Proposition 2. *First Statement:* Total turnout is given by

$$T(m) = \Phi(y_L(m) - m) + 1 - \Phi(y_R(m) - m).$$

Note that as $x_R = -x_L$ and Φ is symmetric around zero, $T(m)$ is symmetric around zero, and hence $T'(0) = 0$.

We now show that $T''(0) < 0$ and hence turnout assumes a local maximum at $m = 0$. Note that

$$T''(0) = (1 - y_L'(0))^2\phi'(y_L(0)) - (1 - y_R'(0))^2\phi'(y_R(0)) + y_L''(0)\phi(y_L(0)) - y_R''(0)\phi(y_R(0)).$$

Differentiating the equilibrium condition $\gamma_L(y_L(m), (-\infty, y_L(m)], S_R) = c$ with respect to m at $m = 0$ with symmetrically-situated candidates, $x_R = -x_L$, yields

$$\Phi(y_L(m))\,\frac{\partial y_L(m)}{\partial m}\Big|_{m=0} + \phi(y_L(m))\left(\frac{\partial y_L(m)}{\partial m}\Big|_{m=0} - 1\right)y_L(m) = 0,$$

which we solve for

$$\frac{\partial y_L(m)}{\partial m} = \frac{y_L(m)\phi(y_L(m) - m)}{\Phi(y_L(m) - m) + y_L(m)\phi(y_L(m) - m)}. \tag{11}$$

Similarly, we get

$$\frac{\partial y_R(m)}{\partial m} = \frac{-y_R(m)\phi(y_R(m) - m)}{1 - \Phi(y_R(m) - m) - y_R(m)\phi(y_R(m) - m)}. \tag{12}$$

In Result 1, we proved that $y_R(m)$ strictly decreases in m. Hence, from (12) we have that the following condition must be satisfied

$$1 - \Phi(y_R(m) - m) - y_R(m)\phi(y_R(m) - m) > 0. \tag{13}$$

Note that the numerators and denominators of (11) and (12) are identical at $m = 0$. Hence,

$$\frac{\partial y_L(m)}{\partial m}\Big|_{m=0} = \frac{\partial y_R(m)}{\partial m}\Big|_{m=0}.$$

Further, the derivatives of the numerators have the same magnitude but opposite signs at $m = 0$. The same is true for the derivatives of the denominators. Thus,

$$\frac{\partial^2 y_L(m)}{\partial m^2}\Big|_{m=0} = -\frac{\partial^2 y_R(m)}{\partial m^2}\Big|_{m=0}.$$

Differentiating (12) with respect to m at $m = 0$ and substituting (12) for the first derivative of y_R yields

$$\frac{\partial^2 y_R(m)}{\partial m^2}\Big|_{m=0} = \frac{(1-\Phi(y_R(0)))y_R(2\phi^2(y_R(0)) + (1-\Phi(y_R(0)))\phi'(y_R(0)))}{(1-\Phi(y_R(0)) - y_R\phi(y_R(0)))^3}. \tag{14}$$

Given that the first derivatives of y_L and y_R are identical, while the second derivatives are of the same magnitude but with opposite signs, we get

$$T''(0) = -2(1 - y_R'^2\phi'(y_R(0)) - 2y_R''(0)\phi(y_R(0)). \tag{15}$$

Substituting the values of the first and second derivatives yields

$$T''(0) = -\frac{2(1-\Phi(y_R(0)))\left(2y_R(0)\phi^3(y_R(0)) + (1-\Phi(y_R(0)))^2\phi'(y_R(0))\right)}{(1-\Phi(y_R(0)) - y_R\phi(y_R(0)))^3}.$$

Note that (13) yields that the denominator is strictly positive. Thus, to show that $T''(0) < 0$ it suffices to show that

$$2y_R(0)\phi^3(y_R(0)) + (1-\Phi(y_R(0)))^2\phi'(y_R(0)) > 0. \tag{16}$$

This is immediate for uniform distribution as $\phi' = 0$.

To provide the proof for normal distributions, note that the sign of (16) is equivalent to that of

$$2y_R(0)\phi^2(y_R(0)) + (1-\Phi(y_R(0)))^2\frac{\phi'(y_R(0))}{\phi(y_R(0))} > 0, \tag{17}$$

which, after substituting for the normal density and its derivative, is equivalent to

$$2y_R(0)\phi^2(y_R(0)) > (1-\Phi(y_R(0)))^2\frac{y_R(0)}{\sigma^2}.$$

Because the normal distribution has a monotone hazard rate and $y_R > 0$, it is sufficient to prove that

$$2\frac{\phi^2(0)}{(1-\Phi(0))^2} > \frac{1}{\sigma^2}.$$

Substituting for the normal density at 0 yields

$$\frac{8}{2\pi\sigma^2} > \frac{1}{\sigma^2},$$

which clearly holds.

Finally, we substitute the cdf of the logistic distribution, $1/(1+e^{-y_R(0)/s})$ for Φ and its derivatives into (17) to obtain:

$$\frac{e^{y_R(0)/s}\left(s - e^{2y_R(0)/s}(s - 2y_R(0))\right)}{(1 + e^{y_R(0)/s})^6 s^3} \tag{18}$$

It is sufficient to prove that $s - e^{2x/s}(s - 2x) > 0$ for $x > 0$ because $y_R(0) > 0$. This expression is 0 at $x = 0$. The derivative with respect to x is $4e^{2x/s}x/s > 0$. Hence, the result follows.

Second Statement: The cases where $c \geq \bar{c}$ follows as above. Thus, let $c < \bar{c}$. Total turnout for $c < \bar{c}$ is

$$T(\eta) = \Phi(y_L(\eta)) + 1 - \Phi(y_R(\eta)).$$

Note that, for $x_R = -x_L$ and given that Φ is symmetric around zero, $T(\eta)$ is symmetric around zero, and hence $T'(0) = 0$.

We now show that $T''(0) < 0$ and hence turnout assumes a local maximum at $\eta = 0$. Note that

$$T''(0) = \phi'(y_L(0))y_L'^2(0) - \phi'(y_R(0))y_R'^2(0) + \phi(y_L(0))y_L''(0) - \phi(y_R(0))y_R''(0)$$

From the proof of Proposition 2, $y_L' = y_R'$, $y_R'' = -y_L'' > 0$. Together with the symmetry of Φ, this gives

$$T''(0) = -2\phi'(y_R(0))y_R'^2(0) - 2\phi(y_R(0))y_R''(0). \tag{19}$$

Note that $y_R'(0) < 0$ implies $y_R'^2(0) < (1 - y_R'^2$. Further, $\phi'(y_R(0)) < 0$. Thus, if (15) is less than 0 then (19) is strictly less than 0. Hence the result holds. ■

9 Online Appendix: Alternative Voting Models

9.1 Abstention by Indifference

In the abstention-by-indifference model, citizens abstain if the difference in utilities from the candidates is below some threshold level, k, i.e., citizen y abstains if $|(x_L - y)^2 - (x_R - y)^2| < k$.

First, a low value of k can generate substantial turnout. Second, if, as one would expect, k is increasing in costs, c, then it is immediate that abstention increases with c. Third, let $S_L = (-\infty, y_L)$ and $S_R = (y_R, \infty)$ be the set of prosocial voters for candidates L and R, respectively. Then total turnout given distribution $\Phi(x-m)$ is

$$T(m) = 1 - (\Phi(y_R - m) - \Phi(y_L - m)). \tag{20}$$

From symmetry, $T'(0) = 0$ and $T''(0) = \phi'(y_L) - \phi'(y_R) = 2\phi'(y_L) > 0$. Thus, turnout obtains a local minimum when the election is most competitive. Moreover, for single peaked distributions turnout continues to increase, and hence the model counterfactually predicts higher expected turnouts in lopsided elections.

There is no bandwagon effect because S_L and S_R are unaffected by perceived changes of the distribution. Finally, it is immediate that abstention is highest for moderates.

9.2 Expressive Voting

When voters have purely selfish preferences ($\beta = 1$), we obtain a closed-form solution for the equilibrium. In particular, $S_L = \left[x_L - \sqrt{-u(N) - c}, x_L + \sqrt{-u(N) - c}\right] \cap (-\infty, \bar{x}]$ if $-u(N) > c$ and $S_L = \emptyset$, otherwise; and similarly for S_R. We now establish the properties of the equilibrium listed in Table 2.

First, note that the turnout rate can be large if $-u(N) - c$ is sufficiently large. Second, by inspection, raising voting costs strictly decreases S_L and S_R whenever they are non-empty. Third, we show how total participation relates to the closeness of the election. Let $S_i = [\underline{y}_i, \bar{y}_i]$. Total turnout given distribution $\Phi(x - m)$ is:

$$T(m) = \Phi(\bar{y}_R - m) - \Phi(\underline{y}_R - m) + \Phi(\bar{y}_L - m) - \Phi(\underline{y}_L - m). \tag{21}$$

It is immediate from symmetry that $T'(0) = 0$. Next,

$$T''(0) = \phi'(\bar{y}_R) - \phi'(\underline{y}_R) + \phi'(\bar{y}_L) - \phi'(\underline{y}_L) = 2\left(\phi'(\bar{y}_R) - \phi'(\underline{y}_R)\right), \tag{22}$$

where the second equality follows from symmetry. The sign is ambiguous, depending on the slopes of ϕ at $\bar{y}$ and $\underline{y}$. Fourth, there is no bandwagon effect because S_L and S_R are unaffected by perceived changes of the distribution. Finally, we consider how abstention relates to ideology. If costs are low enough that a citizen with ideology $\bar{x}$ votes, then abstention is higher only for more extreme types. For higher costs, the abstention pattern is W-shaped.

Finally, for this model to match the empirical facts about lopsided elections, turnout must decrease as the mean of Φ is increased past 0. The argument for abstention by indifference implies that the length of the voting interval must be sufficiently small, else turnout would still increase as elections become more lopsided. This, in turn means, that $\Phi(S_L)$ and/or $\Phi(S_R)$ cannot be very large. To generate turnout that falls in the observed range, there must be substantial turnout by noise voters. This, however, means that the vote margin is largely determined by noise voters, which implies that (i) the expected winning margin is small, and (ii) the underdog wins with non-trivial probability. It follows that we cannot match the empirical regularities of lopsided elections.

9.3 Pivotal Voting Models

The standard pivotal voting model assumes that there is an exogenous source of uncertainty, for example about the number of supporters for each of the two candidates. A citizen votes if the pivot probability, p times the net benefit, B of moving the election from a loss to a tie, or from a tie to win of the favored candidate exceeds the voting costs, c, i.e., if $c < pB$. First, observe that with a large electorate this leads to the paradox of voting, because the pivot probability and hence the turnout percentage converge to zero (see, e.g., Palfrey and Rosenthal (1985)). Second, inspection of the inequality $c < pB$ yields that if c increases then turnout falls. Third, the literature has established that turnout is maximized when the election is expected to be close (see, e.g., Proposition 1 in Myatt (2015)). Fourth, Lemma 4 in Myatt (2015) establishes that there is an underdog effect, i.e., supporters for the underdog turn out at a higher rate than supporters of the leading candidate. Finally, the net benefit B of all citizens with ideology between x_L and x_R is less than those of ideologically more extreme citizens. Thus, abstention is highest among moderates. Finally, by design it is not possible to generate safe elections, as enough uncertainty about the electoral outcome is needed to generate incentives to vote.

9.4 Ethical Voting Models

First, unlike pivotal voting models, ethical voting models can generate substantial participation. Second, increasing voting costs has the standard effect of reducing participation. Third, closer elections result in a higher turnout (Property 3 in Feddersen and Sandroni (2006)). Fourth, as in pivotal voting models, there is an underdog effect (Property 1 in Feddersen and Sandroni (2006)). Standard models in this literature do not consider spatial variation in group composition. To integrate spatial heterogeneity in voter preferences into ethical voting models, we assume that a potential voter for candidate $i = L, R$ with ideology y and voting cost c incurs a cost $u(|y - x_i|) + c$ from voting, where x_i is the candidate's platform and u is increasing in its argument. Further, citizens derive a benefit, $B(y)$ if their preferred candidate wins. Extending the logic of classical ethical voting models, it follows that the set of voters for candidate i should be chosen to maximize group utility. Were ethical behavior solely to revolve around costs, then the set of voters S_L and S_R would be symmetric around x_L and x_R, respectively. In particular, given measures μ_L and μ_R of citizens who should feel obliged to vote, it is immediate that, to maximize group welfare, the sets of citizens that minimize total voting costs are intervals of the form $(x_i - a_i, x_i + a_i)$, $i = L, R$ if $a_i < |x_i|$. Thus, cost minimization implies that the more extreme types should abstain. However, an ethical rule should also account for how much a particular citizen cares about the electoral outcome. For example, if y is very close to $\bar{x} = (x_L + x_R)/2$ then the benefit $B(y)$ becomes negligible, and y should not be ethically compelled to incur voting costs. Hence, there must also be abstention of moderates, resulting in a W-shaped pattern of abstention. Finally, incentives to vote are generated by introducing sufficient noise, e.g., about the realized number of ethical voters. Thus, as in the pivotal voting, an election cannot be truly safe, as the "ethical rule" would want to reduce the number of voters for the safe candidate to a level at which the electoral outcomes become sufficiently uncertain to avoid incurring unneeded voting costs.