Motivated or Mobilized?

Competitiveness, Campaign Effort, and Turnout in U.S. Elections

Luke Miller*

Georgetown University

September 20, 2025

For latest version see here

Abstract

Voter turnout is substantially higher in battleground than in non-battleground states during U.S. presidential elections. It is unclear whether this gap reflects an organic response to closer contests or a manufactured outcome of campaign activity. Existing research treats these channels separately, overlooking the feedback whereby campaign effort shapes voter preferences, alters electoral competitiveness, and in turn affects participation. I develop and estimate a unified structural model in which turnout, competitiveness, and strategic effort are jointly determined in equilibrium. The model is estimated using county-level data from the 2008–2020 presidential elections. Validation exercises show the model's predicted competitiveness effects match reduced-form estimates and that its implied effort allocations align closely with observed advertising patterns. The results indicate that higher competitiveness and strategic effort together fully explain the 6.1 percentage point turnout gap between battleground and non-battleground states, with roughly one-third attributable to competitiveness itself and two-thirds to candidate effort. I use the model to assess the efficiency of this mobilization, finding a marginal cost of about $250 per additional voter, though average costs are much smaller at around $150 per vote. As an application, I simulate a public financing reform that caps campaign budgets, finding that tighter spending limits increase competitiveness but reduce overall turnout.

JEL Classification: D72, C51, C54
Keywords: Voter Turnout, Structural Estimation, Campaign Strategy, Elections, Political Economy

*lm1410@georgetown.edu, https://lukebmiller.github.io/. Many thanks to Laurent Bouton, Juan Felipe Riaño, John Rust, Ami Ko, Allison Stashko, and the participants of several Georgetown Economics seminars for their valuable feedback.

1 Introduction

Voters determine which candidates are elected and, in turn, which policies are enacted. Yet despite their centrality to democratic governance, we still lack a clear account of what motivates citizens to participate in elections. Without such understanding, it is difficult to predict how public preferences translate into electoral outcomes or how institutional or strategic interventions affect turnout. Two standard perspectives dominate. In voter-centric models, closer races increase the instrumental value of voting, encouraging participation. In candidate-centric models, competitiveness induces campaigns to intensify their effort, which shifts voter preferences, again boosting turnout. Both perspectives imply that participation rises with competitiveness, though through distinct channels.

Nowhere is this ambiguity more consequential than in U.S. presidential elections, where turnout is consistently higher in "battleground" states in which margins are narrow and campaign activity is intense (see Table 1). How much of this difference in turnout reflects competitiveness itself, and how much is a product of campaign effort? The distinction matters not only for understanding turnout, but also for policy evaluation. Consider a public financing reform that caps and equalizes spending. Equalized budgets could heighten competitiveness and boost turnout, yet reduced spending may weaken mobilization and depress participation. A credible assessment of such reforms requires a framework that jointly accounts for both channels and quantifies their relative contributions.

This paper develops and estimates a structural model that addresses this challenge. On the voter side, participation depends on costs, partisan leanings, and the competitiveness of the race. On the candidate side, campaigns allocate resources strategically, with effort directly shifting voter preferences and indirectly altering the competitiveness of the race. This unified framework disentangles the roles of competitiveness and campaign effort in driving turnout and provides a basis for counterfactual policy analysis. The model is estimated using county-level data from the 2008–2020 U.S. presidential elections. Although campaign effort is not directly observed, the model infers it as a latent variable consistent with candi-

dates' best responses. The framework also ensures equilibrium consistency: competitiveness governs perceived voting efficacy, which shapes turnout. Turnout, in turn, feeds back into determining the competitiveness of the race.

Table 1: Battleground vs. Non-Battleground States

	Battleground	**Non-Battleground**
Turnout Rate (%)	62.6	56.5
Avg. Margin (%)	2.5	10.2
TV Ad Share (%)	87	13

Notes: Battleground states are defined as the top 10 states by combined Democratic and Republican television advertising expenditures in each presidential election year (2008, 2012, 2016, 2020). Turnout is measured as the share of the voting age population who cast a ballot for the Republican or Democratic candidate in the presidential election.

Identification rests on two elements. First, the economic structure separates effort, which shifts relative candidate utilities directly, from competitiveness, which alters perceived voting efficacy and thereby changes the effective cost of voting. Equilibrium conditions link these components by jointly determining turnout, competitiveness, and candidate effort. Second, exclusion restrictions assign covariates either to partisan preferences (e.g. demographics) or to voting costs (e.g. registration laws), allowing the two channels to be separately identified. Simulations validate this approach, showing that the model can recover the true parameters from synthetic data.

I provide two forms of external validation. First, a border discontinuity design shows the model's predicted turnout response to competitiveness closely matches new quasi-experimental estimates. Second, the model's predicted allocation of effort aligns closely with observed campaign spending patterns.

The model shows that the 6.1 percentage point turnout gap between battleground and non-battleground states is fully explained by a combination of the direct effect of competitiveness on individuals' perceived voting efficacy and the mobilization generated by strategic campaign effort. Decomposition shows that roughly one-third of the gap is attributable to

differences in voting efficacy, while the remaining two-thirds are driven by mobilization.

Having established the core mechanisms, I turn to their implications. First, the model sheds light on the puzzle of why campaigns spend so heavily on mobilization despite small marginal effects in field experiments (Spenkuch and Toniatti, 2018; Aggarwal et al., 2023; Bär, Pröllochs, and Feuerriegel, 2025; Sides, Vavreck, and Warshaw, 2022; Kalla and Broockman, 2018). Early spending mobilizes the most responsive voters. As effort accumulates, marginal returns diminish, but neither candidate can retreat without ceding advantage.

Most reduced-form studies measure spending effects in these saturated environments, where additional outreach has limited impact. The model, by contrast, captures the full equilibrium consequences of strategic escalation, which are quite large, adding nearly three million votes per cycle and raising turnout by 4 to 5 percentage points in pivotal states.

Second, I use the framework to examine a public financing reform that caps campaign budgets at the inflation-adjusted equivalent of the 2008 public grant. The reform highlights a central trade-off: spending caps make races more competitive, raising perceived voting efficacy, but they also reduce mobilization. The net effect is a decline in participation, as lower spending outweighs the effects of intrinsic competitiveness on turnout.

The existing literature has typically examined the roles of electoral competitiveness and campaign effort in isolation. One class of theories focuses on the effects of competitiveness, whereby voters are motivated by the instrumental value of their vote in a close race (Downs, 1957; Palfrey and Rosenthal, 1983). A separate class of theories emphasizes strategic effort, where candidates either mobilize supporters through costly outreach (Herrera, Levine, and Martinelli, 2008) or influence voter choice by targeting resources to undecided voters (Lindbeck and Weibull, 1987; Dixit and Londregan, 1996). This separation has continued into structural models of voter turnout. Candidate-focused models such as Strömberg (2008) and Shachar and Nalebuff (1999) endogenize campaign effort but assume full turnout. Voter-centric models (Coate and Conlin, 2004; Coate, Conlin, and Moro, 2008; Kawai, Toyama, and Watanabe, 2021; Degan and Merlo, 2011) allow turnout to respond to preferences or

voting costs, but treat campaign behavior as exogenous.

However, empirical evidence indicates that both mechanisms matter. Closer contests increase participation (Bursztyn et al., 2024; Levine and Palfrey, 2007; Duffy and Tavits, 2008), while field experiments and observational studies show that campaign effort also raises turnout (Gerber, 2004; Nickerson, 2006; Enos and Fowler, 2018). Models that omit one channel risk misattributing its effects to the other and understate the extent of their interaction. In practice, this means not only that estimates of competitiveness or effort may be misspecified, but also that the dynamic feedback between them is lost: when a trailing candidate intensifies effort, the margin narrows, raising competitiveness and amplifying turnout; when a frontrunner expands spending, the margin widens, dampening competitiveness and reducing the return to further effort.

By integrating voter turnout and campaign strategy within a common framework, this paper bridges two literatures that have traditionally been studied in isolation. The framework clarifies the sources of higher participation in battleground states and provides a tractable tool for evaluating institutional reforms in U.S. presidential elections and beyond.

The paper proceeds as follows. Section 2 presents the model, detailing voter and candidate behavior. Section 3 describes the estimation strategy. Section 4 outlines the data used in the analysis, including campaign spending and turnout measures. Section 5 presents the main results. Section 6 evaluates the model's fit to the data, comparing predicted spending patterns to observed data, as well as the model's predicted turnout responses to changes in competitiveness to reduced-form estimates. Section 7 explores the implications for both voter behavior and campaign strategy. Section 8 summarizes the main findings and outlines directions for future research.

2 The Model

I develop a game-theoretic model of U.S. presidential elections with two sets of actors: voters and candidates, a Democrat (D) and a Republican (R). The candidates are endowed with campaign budgets and allocate resources across states to maximize their probability of winning the Electoral College. This requires securing 270 electoral votes. Each state $s \in \mathcal{S}$ gives its assigned electoral votes l_s to a candidate using a first-past-the-post election. An individual state s is partitioned into counties indexed by $j_s \in \mathcal{J}_s$, where $\sum_{j_s \in \mathcal{J}_s} w_{j_s} = 1$ and w_{j_s} denotes the county's share of the state's voting-eligible population.[1]

Voter and candidate decisions are interdependent. Voters decide whether and for whom to vote based on their partisan preferences, participation costs, and the perceived efficacy of their vote, which is determined by state-level competitiveness. Candidates, in turn, strategically allocate their campaign effort to influence these partisan preferences. The following sections detail the behavior of each actor, beginning with the voter's decision calculus 2.1 before turning to the candidate's equilibrium strategy 2.2.

2.1 Voters

2.1.1 Baseline Preferences

Individuals begin with baseline preferences for the two candidates. Let $\Delta\tilde{u}_{ij_s}$ denote the pre-campaign utility differential between candidate R and candidate D for individual i in county j_s of state s. This is estimated using county-level demographic covariates $X^{\mu}_{j_s}$, as well as unobserved preference shocks. A positive value indicates a preference for R; a negative value indicates a preference for D. Individuals with $\Delta\tilde{u}_{ij_s} > 0$ are classified as aligned with R, and those with $\Delta\tilde{u}_{ij_s} < 0$ as aligned with D. The differential is given by:

[1] For expositional clarity I present the model for a single election year and omit the time index t. In estimation (Section 3.3), the model is applied separately to each presidential election from 2008 to 2020 and the likelihood is pooled across years.

$$\Delta\tilde{u}_{ij_s} = \mu(X^{\mu}_{j_s}) - \eta_{j_s} - \delta_s - \epsilon_{ij_s},$$

where $\mu(X^{\mu}_{j_s})$ captures the average partisan lean of county j_s. For notational convenience, I write μ_{j_s} when the argument is clear. The remaining terms introduce heterogeneity across counties, states, and individuals. Specifically, $\eta_{j_s} \sim \mathcal{N}(0, \sigma_j^2)$ is a county-level deviation from the baseline; $\delta_s \sim \mathcal{N}(0, \sigma_s^2)$ is a common statewide shock; and $\epsilon_{ij_s} \sim \mathcal{N}(0, 1)$ reflects idiosyncratic preference noise at the individual level. The state and county-level shocks capture systematic factors that shift attitudes toward candidate R uniformly across the state or county, such as gubernatorial popularity, salient statewide policy, or national partisan narratives disproportionately affecting a given state or county. All shocks are mean-zero and independently distributed. Uniform national trends, such as presidential approval or macroeconomic sentiment, are absorbed into $\mu(X^{\mu}_{j_s})$. This term captures both local partisan lean and any constant shift in national preferences within a given election cycle.

Positive realizations of the shocks reduce support for R, shifting voters toward D. Because each component enters symmetrically and additively, the county fixed effect $\mu(X^{\mu}_{j_s})$ represents the expected pre-campaign utility advantage of candidate R in county j_s.

2.1.2 Effect of Campaign Effort

Campaigns influence turnout through two mechanisms: they raise the utility differential among aligned voters (mobilization) and reduce it among opponents (persuasion). However, these effects need not be symmetric. For example, campaign rallies may energize supporters while having little effect on the opposing side. To flexibly capture both channels, I define the net utility effect of candidate q and $-q$'s effort as:

$$m(e_{sq}, e_{s-q}) = \theta e_{sq}^{1/\gamma} - \psi\theta e_{s-q}^{1/\gamma}, \tag{1}$$

where $q \in \{D, R\}$ indexes the candidate and $-q$ the opponent. The function is increasing

and concave in own effort and decreasing in the opponent's effort. The scale parameter $\theta > 0$ governs the overall efficacy of campaign effort, $\gamma > 1$ captures diminishing returns, and $\psi \in [0, 1]$ captures the strength of persuasion. This flexible form nests two polar cases: when $\psi = 0$, campaigns affect only aligned voters (pure mobilization); when $\psi = 1$, they exert equal and opposite effects on both sides (full persuasion). The value of e_{sD} and e_{sR} denote the per-capita effort levels and are solved endogenously in equilibrium, as described in Section 2.2.

The overall impact of the campaign environment may differ across counties. Beyond the direct effect of spending, some counties may become more polarized or more moderate as the campaign unfolds. To capture this heterogeneity, I introduce a county-specific shock $\zeta_{j_s} \sim \mathcal{N}(0, \sigma_j^2)$ that uniformly shifts the post-campaign utility differential in county j_s of state s, pushing the electorate toward either stronger partisan alignment or greater moderation, holding effort fixed. This term is distinct from baseline preference shocks but drawn from the same distribution, reflecting a common source of unobserved local heterogeneity.

Post-campaign utility differentials are given by:

$$\Delta u_{ij_s} = \begin{cases} \min\left\{-m(e_{sD}, e_{sR}) + \Delta\tilde{u}_{ij_s} + \zeta_{j_s},\ 0\right\} & \text{if } \Delta\tilde{u}_{ij_s} < 0, \\ \max\left\{m(e_{sR}, e_{sD}) + \Delta\tilde{u}_{ij_s} - \zeta_{j_s},\ 0\right\} & \text{if } \Delta\tilde{u}_{ij_s} \geq 0. \end{cases} \tag{2}$$

So long as $m(e_{sD}, e_{sR}) > 0$, the utility differential moves in favor of the aligned candidate. A positive realization of ζ_{j_s} reduces its absolute magnitude, shifting the electorate toward moderation, while a negative realization increases it, amplifying partisan alignment.

To enforce the model's assumption of fixed partisan alignment, I truncate the post-campaign utility differential so that it cannot cross zero: for voters initially aligned with candidate D (i.e., $\Delta\tilde{u}_{ij_s} < 0$), the post-campaign differential is constrained to remain non-positive, while for voters aligned with R it must remain non-negative. This restriction prevents large effort imbalances or large realizations of ζ_{j_s} from inverting voter types, thereby

maintaining consistency with the fixed-type classification. The truncation is applied to the post-campaign utility differential, not to the individual shock ζ_{j_s}, which remains normally distributed. Substantively, this assumption reflects evidence that campaign activity primarily affects mobilization and demobilization of existing supporters rather than persuading the opposing side to switch allegiances (Schuster (2020)).

2.1.3 Voting Decision and Perceived Efficacy

Voters face a cost of voting, $c(X_{j_s}^c)$, that varies by county. The covariates $X_{j_s}^c$ capture demographic and institutional factors that influence voting costs and are distinct from those affecting baseline preferences. For clarity I write c_{j_s} when the argument is clear. A voter turns out only if the expected benefit of voting exceeds this cost. The expected benefit depends on perceived voting efficacy at the state level, denoted $p(\kappa_s)$.

Rather than literal pivot probabilities, I model perceived voting efficacy as a smooth function of electoral competitiveness. The term *perceived voting efficacy* follows Kawai, Toyama, and Watanabe (2021), who introduce it as a reduced-form object. In contrast, I model efficacy as a structural equilibrium object: it is jointly determined with turnout and competitiveness, and its shape is estimated from the data rather than imposed ex ante.

Competitiveness, κ_s, is defined as the ratio of the losing candidate's vote share to the winner's:

$$\kappa_s \equiv \begin{cases} \dfrac{\sigma_{sD}}{\sigma_{sR}}, & \text{if } \sigma_{sR} > \sigma_{sD}, \\ \dfrac{\sigma_{sR}}{\sigma_{sD}}, & \text{otherwise,} \end{cases} \qquad \kappa_s \in (0, 1]. \tag{3}$$

The terms σ_{sD} and σ_{sR} are population-weighted turnout shares for candidates D and R in state s, respectively:

$$\sigma_{sD} \equiv \sum_{j_s \in \mathcal{J}_s} w_{j_s} \sigma_{j_s D}, \qquad \sigma_{sR} \equiv \sum_{j_s \in \mathcal{J}_s} w_{j_s} \sigma_{j_s R}.$$

I assume that $p(\kappa_s)$ is strictly increasing in electoral competitiveness, reaching its maximum when the election is a tie. In practice, I model $p(\kappa_s)$ as a scaled logistic function:

$$p(\kappa_s) = \frac{\tilde{p}(\kappa_s) - \tilde{p}(0)}{\tilde{p}(1) - \tilde{p}(0)} \qquad \text{where} \quad \tilde{p}(\kappa_s) = \frac{1}{1 + \exp\left(-\alpha_1(\kappa_s - \alpha_2)\right)}.$$

The rescaling normalizes $p(\kappa_s)$ so it remains bounded between 0 and 1. When $\kappa_s = 0$, the election is a landslide and $p(\kappa_s) = 0$; when $\kappa_s = 1$, the election is a perfect tie and $p(\kappa_s) = 1$.

The logistic specification is flexible enough to capture a wide range of turnout responses to competitiveness (Figure 1). In the limit as $\alpha_1 \to \infty$ and $\alpha_2 \to 1$, the function approaches a step, reproducing the classical pivotal-voter logic in which turnout rises only at exact ties. Smaller values of (α_1, α_2) generate smoother, more gradual responses and shift the midpoint, consistent with the idea that voters cannot precisely compute pivot probabilities and instead rely on heuristics about competitiveness. Such comparative statics are consistent with laboratory evidence that individuals tend to overestimate their influence in close elections yet still adjust turnout in response to competitiveness (Levine and Palfrey (2007) and Duffy and Tavits (2008)).

A key feature of my approach is to estimate (α_1, α_2) from the data, allowing the efficacy curve to be disciplined empirically rather than imposed ex ante. In this way, the model nests competing microfoundations of turnout within a common structural framework.

Using the perceived voting efficacy $p(\kappa_s)$, an individual votes if the expected benefit from voting exceeds the cost of voting. The benefit is the product of two components: the utility differential between the two candidates, $|\Delta u_{ij_s}|$, and the perceived efficacy of the vote, $p(\kappa_s)$. The efficacy term, $p(\kappa_s)$ represents a voter's subjective belief in their vote's importance. It serves as a behaviorally plausible and empirically tractable analog to the classical pivot probability derived in the canonical instrumental voter model (see Appendix B.1 for details).

An individual in county j_s votes for D if $\Delta\tilde{u}_{ij_s} < 0$ and:

Figure 1: Perceived voting efficacy as a function of electoral competitiveness

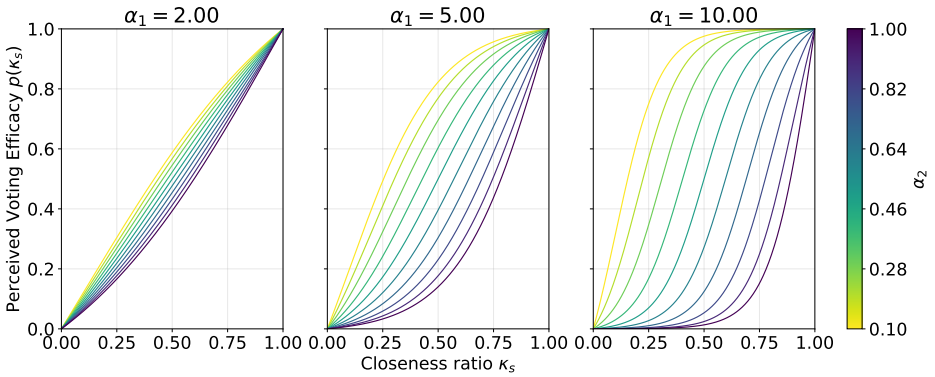

Notes: The efficacy function $p(\kappa_s)$ is a rescaled logistic with steepness α_1 and midpoint α_2. Each panel fixes α_1 at a different level, showing how higher values create a sharper, more step-like response to competitiveness κ_s. Within each panel, colored lines vary α_2, which shifts the midpoint as shown in the color bar.

$$p(\kappa_s) \cdot |\Delta u_{ij_s}| > c_{j_s},$$

and votes for R if $\Delta \tilde{u}_{ij_s} \geq 0$ and:

$$p(\kappa_s) \cdot \Delta u_{ij_s} > c_{j_s}.$$

Otherwise, the individual abstains.

2.1.4 Closed-Form Expressions for Turnout

Let $H(\cdot)$ denote the cumulative distribution function of the idiosyncratic shock ϵ_{ij_s}. The fraction of individuals in county j_s who vote for candidate D is:

$$\sigma_{j_s D} = H\left(m(e_{sD}, e_{sR}) - \mu(X^{\mu}_{j_s}) - \frac{c(X^{c}_{j_s})}{p(\kappa_s)} + \eta_{j_s} + \delta_s - \zeta_{j_s}\right), \tag{4}$$

and the fraction voting for R is:

$$\sigma_{j_sR} = H\left(m(e_{sR}, e_{sD}) + \mu(X^{\mu}_{j_s}) - \frac{c(X^{c}_{j_s})}{p(\kappa_s)} - \eta_{j_s} - \delta_s - \zeta_{j_s}\right). \tag{5}$$

These expressions are derived by finding the threshold for the idiosyncratic shock ϵ_{ij_s} that induces an individual to vote, which requires assuming an interior solution for turnout (i.e., that not all partisans participate). Appendix B.2 provides the detailed derivation.

Finally, let σ_{j_sA} denote the abstention rate in county j_s, with

$$\sigma_{j_sD} + \sigma_{j_sR} + \sigma_{j_sA} = 1.$$

These county-level voting shares and abstention rates depend on perceived voting efficacy, which in turn is a function of state-level turnout aggregates. Because $p(\kappa_s)$ depends on aggregate turnout, and turnout depends on individual voting decisions, equilibrium voting shares are determined through a fixed-point system linking individual behavior and state-level aggregates.

2.1.5 Overview of Voter Decision Mechanism

Figure 2 illustrates the model's voter decision as a three-stage process. First, individuals begin with baseline preferences, represented by the distribution of pre-campaign utility differentials ($\Delta\tilde{u}_{ijs}$). Second, campaign effort and local shocks transform these into post-campaign utility differentials (Δu_{ijs}). Third, individuals decide whether to vote based on their post-campaign utility, the cost of voting (c_{j_s}), and perceived voting efficacy ($p(\kappa_s)$).

The figure illustrates two distinct scenarios that can arise during the second stage. The top row depicts a "Gap Scenario," which occurs when the campaign mobilization effects (m) are larger than the county-specific campaign shock (ζ_{j_s}). This polarizes the electorate, creating an empty space in the post-campaign utility distribution around the indifference point of zero. The bottom row depicts a "Point Mass Scenario," which occurs when a large

shock pushes voters toward indifference. This causes post-campaign utility to be truncated and collected at zero.

Figure 2: Equilibrium turnout and campaign effort within a county

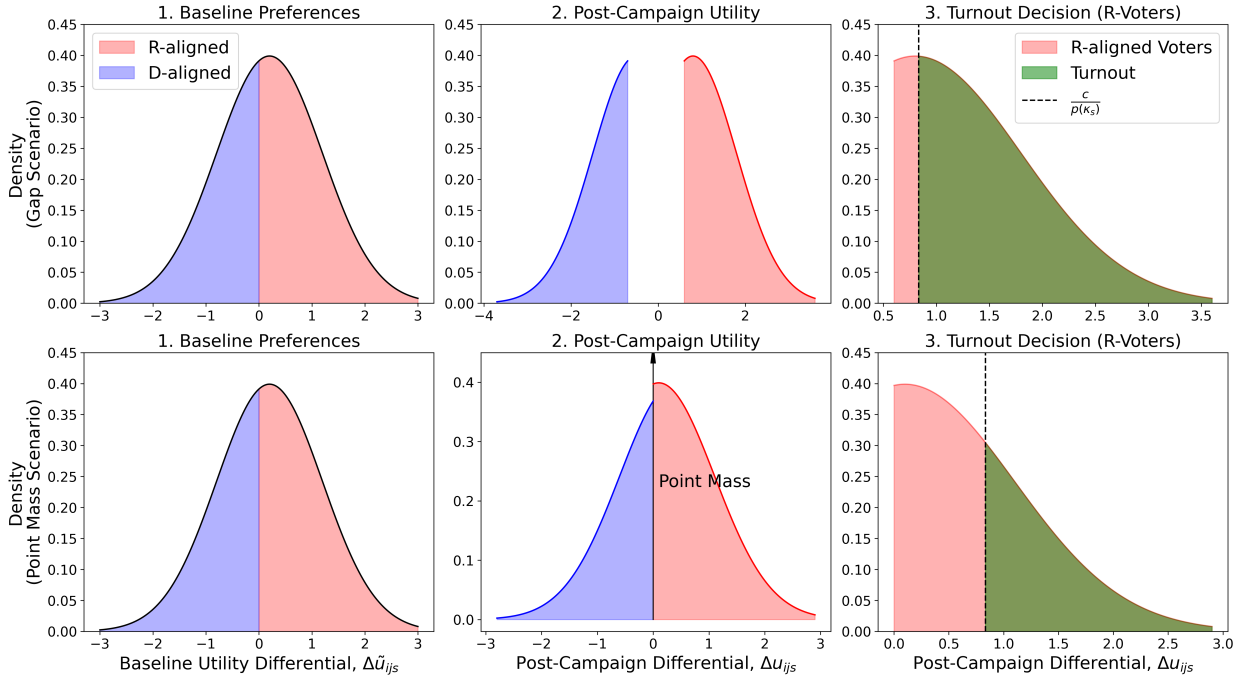

Notes: This figure illustrates the model's three-stage voter decision mechanism under two distinct scenarios. The top row depicts a "Gap Scenario," which occurs when the campaign mobilization effects (m) are larger than the county-specific campaign shock (ζ_{j_s}). This polarizes the electorate, creating an empty space in the post-campaign utility distribution around the indifference point of zero. The bottom row depicts a "Point Mass Scenario," which occurs when a large shock pushes voters toward indifference, where their post-campaign utility is truncated and collects at zero. In both rows, the first panel shows the baseline distribution of pre-campaign utility ($\Delta\tilde{u}_{ijs}$), the second panel shows the transformed post-campaign utility distribution (Δu_{ijs}), and the third panel illustrates the turnout decision for R-aligned voters by showing which fraction of the distribution (shaded green) exceeds the effective cost of voting ($c_{j_s}/p(\kappa_s)$).

2.2 Candidates

2.2.1 Candidates' Objective Function

Each candidate allocates a fixed, exogenously given amount of resources to maximize the probability of winning the election. Let E_D and E_R denote the total campaign resources

available to candidates D and R, respectively.

To reflect the reality of U.S. presidential elections and to reduce computational complexity, I assume that certain states are safe for each party and receive zero campaign effort. Let BG denote the set of battleground states where outcomes are uncertain and candidates allocate positive effort. The safe states contribute their electoral votes automatically to the respective candidate's total. Let EV_D and EV_R denote the number of electoral votes secured by candidates D and R from safe states, respectively.[2]

Candidates concentrate their campaign effort in battleground states and choose allocations to maximize their probability of winning the Electoral College. In making these decisions, they account for how voter turnout responds to campaign effort and perceived electoral competitiveness. Specifically, they consider the distribution of partisan preferences across counties, the cost of voting, and how these factors jointly shape turnout through the perceived efficacy function $p(\kappa_s)$.

While candidates do not observe the realization of local preference shocks η_{j_s}, δ_s, ζ_{j_s}, they are assumed to know their distribution, along with the functional form of turnout behavior. This allows them to compute expected vote shares as a function of both their own and their opponent's campaign effort. Using these expectations, candidates allocate resources across states to maximize the probability that their total electoral vote count meets or exceeds the 270-vote threshold required to win.

Candidate D's problem can be written as:

$$\max_{\{e_{sD}\}_{s\in BG}} \Pr\left(EV_D + \sum_{s\in BG} D_s l_s \geq 270\right) \quad \text{subject to} \quad \sum_{s\in BG} n_s \cdot e_{sD} \leq E_D,$$

where n_s is the voting age population in state s, l_s denotes the number of electoral votes

[2]For empirical support for the battleground / non-battleground state distinction, see, e.g., `https://fairvote.org/report/2008-s-shrinking-battleground-and-its-stark-impact-on-campaign-activity/` and author calculations based on data from the Wesleyan Media Project, which show that over 87% of television advertising and 90% of campaign visits during the 2008–2020 presidential elections occurred in just 10 states on average.

in state s, and $D_s = 1$ if candidate D wins state s, and zero otherwise.[3]

The probability that candidate D wins state s, denoted $\pi_s(e_{sD}, e_{sR})$, is the probability that D's aggregate vote share exceeds that of R, conditional on the candidates' effort levels:

$$\pi_s(e_{sD}, e_{sR}) \equiv \Pr\left(\sigma_{sD}(e_{sD}, e_{sR}) > \sigma_{sR}(e_{sD}, e_{sR})\right).$$

Here, $\sigma_{sD}(e_{sD}, e_{sR})$ and $\sigma_{sR}(e_{sD}, e_{sR})$ denote the state-level turnout shares for D and R, respectively, defined as population-weighted averages of county-level turnout. From each candidate's perspective, these aggregates are random variables: they depend on campaign effort as well as unobserved preference shocks δ_s, η_{j_s}, and ζ_{j_s}, whose distributions are known but whose realizations are not.

Computing the exact distribution of total electoral votes is intractable due to the high-dimensional integration over state and county-specific shocks. To overcome this, I approximate the distribution of the total electoral votes obtained by candidate D using the Central Limit Theorem.[4] Specifically, $\sum_{s \in BG} D_s l_s$ is approximated by a normal distribution with mean and variance:

$$\sum_{s \in BG} D_s l_s \sim \mathcal{N}\left(\sum_{s \in BG} \pi_s(e_{sD}, e_{sR}) l_s, \sum_{s \in BG} \pi_s(e_{sD}, e_{sR})(1 - \pi_s(e_{sD}, e_{sR})) l_s^2\right).$$

Under this approximation, candidate D's objective function becomes:

$$\max_{\{e_{sD}\}_{s \in BG}} \quad \Phi\left(\frac{\sum_{s \in BG} \pi_s(e_{sD}, e_{sR}) \cdot l_s - (270 - EV_D)}{\sqrt{\sum_{s \in BG} \pi_s(e_{sD}, e_{sR}) \left[1 - \pi_s(e_{sD}, e_{sR})\right] l_s^2}}\right) \tag{6}$$

$$\text{subject to} \quad \sum_{s \in BG} n_s \cdot e_{sD} \leq E_D$$

[3] I abstract from states like Maine and Nebraska, which allocate electoral votes by congressional district.

[4] Strömberg (2008) follows a similar approach in approximating expected electoral vote counts under conditional independence across states. Simulations in Appendix B.4 confirm that this approximation is accurate even when uncertainty is limited to battleground states.

where $\Phi(\cdot)$ denotes the cumulative distribution function of the standard normal distribution.

Similarly, candidate R's objective function is:

$$\max_{\{e_{sR}\}_{s\in BG}} \quad \Phi\left(\frac{\sum_{s\in BG}\left[1-\pi_s(e_{sD},e_{sR})\right]\cdot l_s-(270-EV_R)}{\sqrt{\sum_{s\in BG}\pi_s(e_{sD},e_{sR})\left[1-\pi_s(e_{sD},e_{sR})\right]l_s^2}}\right) \tag{7}$$
$$\text{subject to} \quad \sum_{s\in BG} n_s\cdot e_{sR}\leq E_R$$

2.2.2 Calculating the Probability of Winning a State

The probability a candidate wins a state depends on individual-level voting decisions. For candidate D, the probability of winning state s is:

$$\pi_s(e_{sD},e_{sR})=\Pr\left(\sum_{j_s} w_{j_s}\sigma_{j_sD}(e_{sD},e_{sR})-\sum_{j_s} w_{j_s}\sigma_{j_sR}(e_{sD},e_{sR})>0\right) \tag{8}$$

where, following the structure from Section 2.1.4,

$$\sigma_{j_sD}(e_{sD},e_{sR})=H\left(m(e_{sD},e_{sR})-\mu_{j_s}-\frac{c_{j_s}}{p(\kappa_s)}+\eta_{j_s}+\delta_s-\zeta_{j_s}\right),$$
$$\sigma_{j_sR}(e_{sD},e_{sR})=H\left(m(e_{sR},e_{sD})+\mu_{j_s}-\frac{c_{j_s}}{p(\kappa_s)}-\eta_{j_s}-\delta_s-\zeta_{j_s}\right).$$

There are two main challenges in directly solving for the probability in (8). First, the probability depends on three random components: the county-specific shocks η_{j_s} and ζ_{j_s}, and the state-specific shock δ_s. Second, the county-level turnout rates σ_{j_sD} and σ_{j_sR} depend endogenously on the statewide aggregates σ_{sD} and σ_{sR}.

To simplify, I approximate the probability of winning a state by assuming that candidates treat county-level shocks η_{j_s} and ζ_{j_s} as negligible relative to the state-level shock δ_s. This is justified when the variance of county-level shocks is small compared to the variance of

the state-level shock and there are many counties in the state, as shown in Appendix B.5. Fortunately, the recovered parameters in Section 5 match this condition.

Under this approximation, define the adjusted county-level turnout shares as:

$$\tilde{\sigma}_{j_sD}(e_{sD}, e_{sR}) \equiv H\left(m(e_{sD}, e_{sR}) - \mu_{j_s} - \frac{c_{j_s}}{p(\tilde{\kappa}_s)} + \delta_s\right),$$

$$\tilde{\sigma}_{j_sR}(e_{sD}, e_{sR}) \equiv H\left(m(e_{sR}, e_{sD}) + \mu_{j_s} - \frac{c_{j_s}}{p(\tilde{\kappa}_s)} - \delta_s\right),$$

where $\tilde{\kappa}_s$ is identical to equation (3) but uses the adjusted state-level turnout shares calculated as $\tilde{\sigma}_{sD} \equiv \sum_{j_s} w_{j_s}\tilde{\sigma}_{j_sD}$ and $\tilde{\sigma}_{sR} \equiv \sum_{j_s} w_{j_s}\tilde{\sigma}_{j_sR}$.

The probability that the Democratic candidate wins state s is then approximated by:

$$\pi_s(e_{sD}, e_{sR}) \approx \Pr\left(\sum_{j_s} w_{j_s}\tilde{\sigma}_{j_sD}(e_{sD}, e_{sR}) - \sum_{j_s} w_{j_s}\tilde{\sigma}_{j_sR}(e_{sD}, e_{sR}) > 0\right).$$

Let $\hat{\delta}_s$ be the threshold value of δ_s at which the weighted support for D and R exactly balance, following the threshold-based approach in Bouton et al. (2023):

$$\sum_{j_s} w_{j_s} H\left(m(e_{sD}, e_{sR}) - \mu_{j_s} - \frac{c_{j_s}}{p(\tilde{\kappa}_s)} + \hat{\delta}_s\right) =$$
$$\sum_{j_s} w_{j_s} H\left(m(e_{sR}, e_{sD}) + \mu_{j_s} - \frac{c_{j_s}}{p(\tilde{\kappa}_s)} - \hat{\delta}_s\right). \tag{9}$$

When $\delta_s > \hat{\delta}_s$, candidate D wins the state; otherwise, candidate R wins. Thus, conditional on effort profiles (e_{sD}, e_{sR}), the probability that candidate D wins state s is:

$$\pi_s(e_{sD}, e_{sR}) \approx 1 - F(\hat{\delta}_s) \equiv \tilde{\pi}_s(e_{sD}, e_{sR}), \tag{10}$$

where $F(\cdot)$ denotes the cumulative distribution function of the state-level shock δ_s. Substituting $\tilde{\pi}_s(e_{sD}, e_{sR})$ for $\pi_s(e_{sD}, e_{sR})$ yields the following objective functions:

For candidate D:

$$\max_{\{e_{sD}\}_{s\in BG}} \Phi\left(\frac{\sum_{s\in BG}\tilde{\pi}_s(e_{sD},e_{sR})l_s-(270-EV_D)}{\sqrt{\sum_{s\in BG}\tilde{\pi}_s(e_{sD},e_{sR})(1-\tilde{\pi}_s(e_{sD},e_{sR}))l_s^2}}\right) \tag{11}$$
$$\text{subject to} \quad \sum_{s\in BG} n_s\cdot e_{sD}\le E_D,$$

and for candidate R:

$$\max_{\{e_{sR}\}_{s\in BG}} \Phi\left(\frac{\sum_{s\in BG}(1-\tilde{\pi}_s(e_{sD},e_{sR}))l_s-(270-EV_R)}{\sqrt{\sum_{s\in BG}\tilde{\pi}_s(e_{sD},e_{sR})(1-\tilde{\pi}_s(e_{sD},e_{sR}))l_s^2}}\right) \tag{12}$$
$$\text{subject to} \quad \sum_{s\in BG} n_s\cdot e_{sR}\le E_R.$$

2.2.3 Equilibrium Strategies

Given the approximated win probabilities $\tilde{\pi}_s(e_{sD}, e_{sR})$, I solve for a pure-strategy Nash equilibrium in campaign effort. Each candidate $q \in \{D, R\}$ chooses an effort allocation $\{e^*_{s,q}\}_{s\in BG}$ that maximizes their probability of winning the Electoral College, taking the opponent's allocation as given (see equations (11) and (12)).

Importantly, each candidate internalizes how their effort affects turnout via the perceived voting efficacy channel, and thus indirectly affects the probability of winning each state. Equilibrium conditions correspond to the Karush-Kuhn-Tucker (KKT) system associated with each candidate's constrained optimization problem. I solve this system jointly for both candidates using an iterative root-finding algorithm, updating effort vectors until convergence to a fixed point at which neither candidate has a profitable deviation. Simulations that hold county characteristics and structural parameters fixed, and repeatedly vary the initial effort allocations, consistently converge to the same solution, indicating that the equilibrium allocation is unique for a given set of budget constraints and electoral environment.

2.3 Model Equilibrium and Timing

The model's equilibrium is defined by the interaction between candidates' strategic effort choices and the resulting voter turnout. It consists of two interconnected stages.

First, observing baseline partisan preferences and county-level voting costs, candidates choose their effort allocations $\{e^*_{s,D}, e^*_{s,R}\}_{s \in BG}$ in a pure-strategy Nash equilibrium. Because the relevant shocks are not yet realized, candidates maximize their expected probability of winning the Electoral College, taking expectations over the known distribution of the state-level shock δ_s. In deciding their effort allocations, candidates take into account how their efforts shape both the intensity of voter preferences and the perceived efficacy of voting, for any given realization of δ_s.

Once candidate strategies are set and the state- and county-level shocks are realized, a voter-turnout equilibrium is reached. This equilibrium is a fixed point in state-level competitiveness κ_s: voters' turnout decisions, shaped by their perceived efficacy $p(\kappa_s)$, generate aggregate vote shares that reproduce an identical level of competitiveness κ_s.

3 Estimation

3.1 Estimation Framework

I estimate the model using county-level vote shares from U.S. presidential elections between 2008 and 2020. The parameters of interest are a vector of coefficients,

$$\beta \;=\; \big(\beta_\mu,\, \beta_c,\, \beta_{\alpha_1},\, \beta_{\alpha_2},\, \beta_\theta,\, \beta_j,\, \beta_s\big),$$

which govern county-level heterogeneity in partisanship and voting costs, the curvature of the pivotality function, the responsiveness of turnout to campaign effort, and the variances of county- and state-level shocks. All variances are normalized relative to the individual-level shocks ϵ_i, which are assumed to follow a standard normal distribution. Two additional pa-

rameters, γ, the inverse elasticity of campaign production, and ψ, the persuasion parameter, are estimated by grid search.

For each parameter vector β, the model solution yields equilibrium effort allocations and state-level closeness. These inputs allow inversion of the turnout equations to obtain the implied preference shocks. Let $\hat{\eta}_{j_s t}$, $\hat{\zeta}_{j_s t}$, and $\hat{\delta}_{st}$ be the inferred county- and state-level shocks at time t. The likelihood is then maximized over β:

$$\hat{\beta} \;=\; \arg\max_{\beta} \sum_{t=2008}^{2020} \sum_{s \in \mathcal{S}} \sum_{j_s \in \mathcal{J}_s} \log \phi\Big(\hat{\eta}_{j_s t},\, \hat{\zeta}_{j_s t},\, \hat{\delta}_{st} \;\big|\; \beta;\, X^{\mu}_{j_s t},\, X^{c}_{j_s t}\Big),$$

where $\phi(\cdot)$ is the joint density of the inferred shocks under the model's Gaussian assumptions.

When evaluating this likelihood function, the effort terms (e_{sD}, e_{sR}) are treated differently depending on the state's classification. For states identified as battlegrounds (see section 4.4), effort is the endogenously solved equilibrium value. For all non-battleground states, effort is set to zero, consistent with the model's simplifying assumption.

A central advantage of the framework is that state-level effort does not enter as a measured regressor; instead, effort is solved in equilibrium subject to exogenous party budgets. This mitigates concerns about endogeneity in reduced-form designs. The remainder of this section details the parameters of interest, the estimation procedure, and validation of the model's identification.

3.2 Parameters

Table 2 summarizes the structural parameters and shows how each is recovered from the estimated coefficient vector β. County-level covariates $X^{\mu}_{j_s t}$ and $X^{c}_{j_s t}$ capture partisanship fundamentals (e.g. race, education) and voting costs (e.g. polling-place congestion, mail-in availability), respectively. Applying the exponential function to c_{j_s}, θ, and the shock variances ensures they are positive. For the voting-efficacy midpoint α_2, I use the logistic function so that $\alpha_2 \in (0, 1)$. The campaign-effort parameters θ, γ, and ψ jointly govern how

spending translates into voter utility. I estimate θ via maximum likelihood, while γ and ψ are selected by grid search. Preliminary simulations indicated that jointly estimating all three parameters led to local minima and weak identification, particularly for ψ. To ensure stable estimation, I fix γ and ψ over a discrete grid and estimate θ for each combination, selecting the pair that yields the highest likelihood. The final specification uses $\gamma = 2$ and $\psi = 0.25$.

Table 2: Structural parameters and their link to estimated coefficients

Parameter	Economic interpretation	Mapping from β
Partisanship and cost		
$\mu_{j_s t}$	Baseline utility advantage of R in county j_s	$\beta_{\mu t}^{\top} X_{j_s t}^{\mu}$
c_{j_s}	Cost of voting in county j_s	$\exp\big(\beta_c^{\top} X_{j_s t}^{c}\big)$
Perceived voting efficacy		
α_1	Logistic parameter (steepness)	$\exp(\beta_{\alpha_1})$
α_2	Logistic parameter (midpoint)	$1/(1+\exp(-\beta_{\alpha_2}))$
Campaign-effort technology		
θ	Scale of utility gain per \$ spent	$\exp(\beta_\theta)$
γ	Concavity of campaign production function	Grid search: $\gamma \in \{2, 3, 4, 5\}$
ψ	Persuasion parameter	Grid search: $\psi \in \{0.0, 0.1, 0.25, 0.5, 0.75, 1.0\}$
Aggregate shocks		
σ_j	Std. dev. of county-level shock	$\exp(\beta_j)$
σ_s	Std. dev. of state-level shock	$\exp(\beta_s)$

Notes: All exponentiated mappings impose positivity. Covariates in $X_{j_s,t}^{\mu}$ include demographic composition. Covariates in $X_{j_s,t}^{c}$ include polling-place congestion, mail-in options, voter ID laws, median income, and employment rate. The campaign-effort parameters γ and ψ are selected by grid search, while all other coefficients are estimated by maximum likelihood. The partisanship coefficients $\beta_{\mu t}$ are estimated separately for each election year. All other parameters are estimated using pooled data across all years.

Importantly, I assume that the behavioral parameters $(\alpha_1, \alpha_2, \theta, \gamma, \psi)$ and the shock variances (σ_j, σ_s) are constant across election cycles. Likewise, the cost-function coefficients β_c are applied uniformly across years. In contrast, the partisanship component $\mu_{j_s t}$ is modeled as $\mu_{j_s t} = \beta_{\mu t}^{\top} X_{j_s t}^{\mu}$, with both the covariates $X_{j_s t}^{\mu}$ and the coefficients $\beta_{\mu t}$ allowed to vary by election year. This structure captures evolving patterns of local partisanship while holding the underlying behavioral model fixed.

3.3 Estimation Procedure

The structural parameters are estimated by maximum likelihood.[5] The estimation procedure inverts the model's turnout equations to recover the unobserved preference shocks $(\eta_{j_s}, \zeta_{j_s}, \delta_s)$ that rationalize the observed vote shares $(\sigma_{j_sD}, \sigma_{j_sR})$ in equilibrium.

For county j in state s, under an interior solution, the observed Democratic and Republican turnout shares satisfy:

$$\sigma_{j_sD} = H\Big(m(e_{sD}, e_{sR}) - \mu_{j_s} - \tfrac{c_{j_s}}{p(\kappa_s)} + \eta_{j_s} + \delta_s - \zeta_{j_s}\Big),$$
$$\sigma_{j_sR} = H\Big(m(e_{sR}, e_{sD}) + \mu_{j_s} - \tfrac{c_{j_s}}{p(\kappa_s)} - \eta_{j_s} - \delta_s - \zeta_{j_s}\Big),$$

where $H(\cdot)$ denotes the standard normal CDF. Inverting the CDF isolates the latent thresholds:

$$H^{-1}(\sigma_{j_sD}) = m(e_{sD}, e_{sR}) - \mu_{j_s} - \tfrac{c_{j_s}}{p(\kappa_s)} + \eta_{j_s} + \delta_s - \zeta_{j_s}, \tag{13}$$
$$H^{-1}(\sigma_{j_sR}) = m(e_{sR}, e_{sD}) + \mu_{j_s} - \tfrac{c_{j_s}}{p(\kappa_s)} - \eta_{j_s} - \delta_s - \zeta_{j_s}. \tag{14}$$

Adding and subtracting (13) and (14) yields closed-form expressions for ζ_{j_s} and the combined shock $\eta_{j_s} + \delta_s$:

$$\zeta_{j_s} = \tfrac{1}{2}\Big[-H^{-1}(\sigma_{j_sD}) - H^{-1}(\sigma_{j_sR}) + m(e_{sD}, e_{sR}) + m(e_{sR}, e_{sD}) - 2\tfrac{c_{j_s}}{p(\kappa_s)}\Big], \tag{15}$$
$$\eta_{j_s} + \delta_s = \tfrac{1}{2}\Big[H^{-1}(\sigma_{j_sD}) - H^{-1}(\sigma_{j_sR}) - m(e_{sD}, e_{sR}) + m(e_{sR}, e_{sD}) + 2\mu_{j_s}\Big]. \tag{16}$$

Since η_{j_s} is i.i.d. with mean zero, averaging (16) across counties in state s yields an

[5] Throughout I suppress the time index t and treat a single election as given. Coefficients for μ are estimated separately by year, allowing the influence of demographic factors to vary across elections. Coefficients for c are held fixed over time, though the covariates themselves (e.g. polling-place density, voting rules) may change across years. The structural parameters $(\alpha_1, \alpha_2, \theta, \gamma, \psi, \sigma_j, \sigma_s)$ are estimated using pooled data across all years.

unbiased estimator for the state-level preference shock:

$$\hat{\delta}_s = \frac{1}{|J_s|} \sum_{j_s} \tfrac{1}{2} \Big[H^{-1}(\sigma_{j_s D}) - H^{-1}(\sigma_{j_s R}) - m(e_{sD}, e_{sR}) + m(e_{sR}, e_{sD}) + 2\mu_{j_s} \Big]. \tag{17}$$

Subtracting $\hat{\delta}_s$ from (16) recovers the county-specific deviation $\hat{\eta}_{j_s}$.

The direct inversion is valid so long as the observed vote shares for D (R) is less than the fraction of voters aligned with D (R) in the county, as well as the assumption that all observed vote shares lie strictly between 0 and 1.[6] For any such share, the inverse CDF $H^{-1}(\cdot)$ is finite and uniquely defined, allowing one to recover the latent thresholds in (13)–(14). Using the estimated equilibrium efforts (e_{sD}, e_{sR}) these thresholds yield unique estimates of $\hat{\eta}_{j_s}$, $\hat{\zeta}_{j_s}$, and $\hat{\delta}_s$ via (15), (16), and (17).

For a given coefficient vector β, grid values (γ, ψ), and observed vote shares $(\sigma_{j_s D}, \sigma_{j_s R})$, the implied shocks $(\hat{\eta}_{j_s}, \hat{\zeta}_{j_s}, \hat{\delta}_s)$ are uniquely determined. Under the maintained assumption that $\eta_{j_s}, \zeta_{j_s} \sim \mathcal{N}(0, \sigma_j^2)$ and $\delta_s \sim \mathcal{N}(0, \sigma_s^2)$, the log-likelihood equals the sum of Gaussian log-densities evaluated at these inferred shocks.

For each trial β, I solve the inner-outer fixed-point system described in Sections 2.2.2 and 2.2.3 to obtain the equilibrium effort profiles of the candidates. Using these efforts, along with observed county demographics, voting-cost features, and observed vote shares, I recover the implied shocks and evaluate the Gaussian log-likelihood. This derivation illustrates the inversion for a single election year. The full estimation maximizes the pooled likelihood across the 2008–2020 elections, repeating the inversion for each t and aggregating the log-densities over time.

3.4 Identification

Identification of the structural parameters rests on two features of the model and the data.

[6] A preliminary check shows that this is violated in seven counties where the number of registered voters is less than the number of votes cast, likely due to data errors or misreporting. Given this small number of cases, these counties will be excluded from the final estimation.

First, the economic structure distinguishes how competitiveness and effort enter the turnout equations. Campaign effort shifts relative candidate utilities directly through $m(e_{sD}, e_{sR})$, while competitiveness enters only through the perceived efficacy function $p(\kappa_s)$, which scales voting costs. This separation allows the efficacy parameters (α_1, α_2) to be identified from variation in turnout responses to competitiveness, while the effect of campaign effort is identified from how turnout responds to changes in the distribution of resources across states. Campaign effort itself is determined endogenously in equilibrium, linking it to competitiveness and turnout in a manner consistent with candidates' strategic incentives.

Second, exclusion restrictions assign covariates either to baseline partisan preferences, $\mu_t(X_{js}^{\mu})$, or to voting costs, $c(X_{js}^{c})$. Demographic and socioeconomic variables such as race and education shift partisan utilities, while institutional and economic variables such as polling-place congestion or registration deadlines shift participation costs. This variation separates baseline political alignment from participation costs, providing leverage to distinguish the two channels.

County-level heterogeneity plays a central role. Counties within the same state share the same competitiveness and campaign environment, but differ in demographics and voting costs. Comparing turnout across such counties helps identify the effects of $\mu_t(X_{js}^{\mu})$ and $c(X_{js}^{c})$ separately from the common state-level efficacy term $p(\kappa_s)$ and campaign effort. Across states, variation in competitiveness κ_s and campaign effort allows for identification of the efficacy function parameters (α_1, α_2) and the campaign-effectiveness parameters (θ, γ, ψ).

I validate this approach in two ways: (i) by using simulations to show that the estimator can recover known parameters from synthetic data generated by the model, and (ii) by inspecting the likelihood surface around the estimated parameters to confirm that it is locally well-behaved and concave. For more information, see Appendix E.1 and Appendix E.2, respectively. Together, these exercises provide strong support for the model's identification.

Finally, it is crucial to distinguish the identification of the model's structural parameters from the estimation of coefficients on observable covariates. Covariates enter the model

through the relationships $\mu_{jst} = \beta_{\mu t}^{\top} X_{jst}^{\mu}$ and $c_{jst} = \exp(\beta_c^{\top} X_{jst}^{c})$. Their purpose is not to yield causal estimates, but rather to allow baseline partisanship and voting costs to vary flexibly with county characteristics. Accordingly, the coefficients $(\beta_{\mu t}, \beta_c)$ are treated as descriptive summaries of these relationships. The central analytical goal is the identification of the core structural parameters $(\alpha_1, \alpha_2, \theta, \gamma, \psi, \sigma_j, \sigma_s)$ that govern voter and campaign strategy. The identification of these unobserved parameters relies on the model's equilibrium structure.

4 Data

The empirical analysis combines county-election information for the 2008, 2012, 2016, and 2020 presidential cycles. Five public sources are merged to obtain (i) official vote totals, (ii) demographic data, (iii) proxies for the cost of voting, and (iv) campaign budgets. The resulting panel contains $12,225$ county-election observations after minimal sample restrictions described below.

4.1 Voting Data

County-level Democratic and Republican vote counts are taken from the MIT Election Data and Science Lab's harmonized returns. Table 3 summarizes turnout rates by party and geography, disaggregating between county- and state-level aggregates and between battleground and non-battleground states. Two patterns emerge.

First, turnout varies much more across counties than across states. Within battleground states, the standard deviation of county-level turnout is roughly three times larger than the state-level standard deviation. A similar pattern holds in non-battleground states, though the difference is smaller. This dispersion reflects substantial within-state heterogeneity in partisanship and voting conditions, which motivates the use of county-level data for identification. Because competitiveness and campaign effort are determined at the state level, counties within a state share a common strategic environment. The model exploits this

structure by comparing turnout across counties with identical competitiveness and effort, but differing in demographics and voting costs. This variation helps separately identify the effects of baseline utilities, cost shifters, and unobserved shocks on turnout behavior.

Second, state-level turnout is substantially more dispersed in non-battleground states than in battlegrounds, with the standard deviation nearly three times larger in the former. This pattern is consistent with the model's structure, in which turnout depends on effective voting costs $c_{j_s}/p(\kappa_s)$. In battleground states, where competitiveness κ_s is uniformly high, counties face similar levels of voting efficacy $p(\kappa_s)$, resulting in more compressed turnout. In non-battlegrounds, heterogeneity in κ_s introduces variation in $p(\kappa_s)$, which raises dispersion in effective costs and thus in turnout. Although differences in partisan preferences μ_{j_s} or baseline costs c_{j_s} could also contribute, later results show these covariates are similar across both groups, supporting the interpretation that variation in efficacy is a key source of turnout dispersion across states.

Table 3: Turnout Summary Statistics by Geography and State Type

		Battleground States		**Non-Battleground States**	
Level	Party	Mean	Std. Dev.	Mean	Std. Dev.
County	Democrat	0.260	0.093	0.203	0.102
	Republican	0.362	0.097	0.369	0.117
State	Democrat	0.321	0.038	0.262	0.098
	Republican	0.307	0.027	0.275	0.092

Notes: County-level figures are unweighted averages across counties; state-level figures reflect population weighted averages across counties within each state. Turnout is defined as the fraction of eligible voters casting a ballot for each party.

4.2 Demographics Data

Demographic covariates come from the five-year American Community Survey (ACS). For each election t I use the ACS file whose midpoint year equals t.[7] I additionally calculate the density of the county by taking the land area in kilometers squared and dividing it by the number of residents. I use this to classify counties as urban or rural, with the cutoff set at 350 residents per square kilometer.[8] I drop any county-election cell with missing vote totals or demographic information. Further, I drop counties with 500 or fewer residents, as often these areas report more than 100 percent turnout due to small numbers of votes and total population. Tables 21 and 22 in the appendix present summary statistics for the demographic covariates used in the analysis, split by battleground and non-battleground states. The demographic covariates are similar across battleground and non-battleground states.

4.3 Cost of Voting

I use three sets of covariates to capture the cost of voting: (1) a measure of polling-place congestion, (2) state-level election-law cost shifters, and (3) county-level employment rates.

To proxy for queues at the polls, for each county I construct a congestion index that divides the total voting-age population by the number of in-person polling locations on Election Day. This index captures the relative burden of waiting in line to vote, as a higher value indicates more people share a polling place, which can lead to longer wait times and thus increased costs of voting.

Polling place data come from the Election Administration and Voting Survey (EAVS), a biennial national survey administered by the U.S. Election Assistance Commission. The EAVS collects detailed administrative data from local election officials, including the number

[7] For example, the 2006-10 ACS file is matched to the 2008 election.

[8] This threshold is chosen to match the classification rule used by the U.S. Census Bureau which used 1,000 residents per square mile. https://www.census.gov/newsroom/blogs/random-samplings/2022/12/redefining-urban-areas-following-2020-census.html

of polling places, voting methods, and ballots cast. However, approximately 19% of county-election observations are missing polling place data. Appendix C.2 provides details on the imputation strategy for these missing values.

To match the scale of other covariates, I convert this index to its empirical percentile. Motivated by evidence that urban and rural residents experience polling-places differently (Bagwe, Margitic, and Stashko, 2022), I interact the congestion index with an indicator for urban status, defined as counties with more than 350 residents per square kilometer. This yields two separate covariates: one for urban congestion and one for rural congestion. Finally, in states that conduct universal vote-by-mail, I set the congestion index to zero.

To capture cross-state and temporal variation in the administrative burden of voting, I incorporate three state-level measures from Li, Pomante, and Schraufnagel (2018). First, voter identification (ID) requirements are coded on a five-point scale ranging from 0 (signature only) to 4 (strictly enforced photo ID), with intermediate values for less stringent enforcement and non-photo alternatives. Second, polling access is proxied by the statutory number of in-person voting hours on Election Day. Li, Pomante, and Schraufnagel (2018) compute each state's average from the minimum and maximum legal polling times and subtract this value from 20, which is the maximum feasible window from midnight to 8:00 p.m. to define a poll-hours gap. For mail-in voting states, HI, UT, and CO are scored based on actual polling hours since they maintain physical polling places, while OR and WA are assigned the full 20-hour window based on ballot return deadlines. Third, registration deadlines are measured as the number of days before Election Day by which voters must register; states offering same-day registration receive a score of zero.

In all cases, I apply a linear transformation subtracting the minimum and dividing by the range to map the raw values from Li, Pomante, and Schraufnagel (2018) to the unit interval $[0, 1]$, where higher scores correspond to more restrictive policies. This transformation ensures that these election-law cost shifters are on the same scale as the demographic and congestion covariates, making them easier to estimate alongside the other parameters in the

model.

Finally, to account for the opportunity cost of time and broader economic constraints, I include the share of the voting-age population that is employed. Employment shares come directly from American Community Survey data.

Summary statistics for all three sets of cost covariates are presented in Appendix C.1.

4.4 Identifying Battleground States and Constructing Campaign Effort

While the model endogenously determines how effort is allocated across states, I must first identify which states are contested and quantify the total campaign resources available to each party in each election year. I define battleground states empirically using data from the Wesleyan Media Project, which reports state-level television advertising by party and sponsor type. I then use national disbursement records from the Federal Election Commission (FEC) to scale these media expenditures into a broader measure of total mobilization effort within each battleground state. Summing across these states yields the total campaign budget available to each party in each election year.

Battleground states are classified as the ten states with the highest combined Democratic and Republican television advertising expenditures between August 1 and Election Day in each presidential cycle, a timeframe chosen to exclude primary-season spending. Campaign activity is highly concentrated: these states account for 87% of total advertising expenditures, while most of the remaining states receive less than 0.5% individually. Table 12 shows how the spending share rises as states are added in descending order of advertising expenditures: the first few additions yield large gains, but by the tenth state further increases are negligible, indicating that the top ten capture nearly all campaign resources. As a robustness check, I re-estimate the model using the top twelve states; the results are qualitatively similar.

Although a few non-battleground states receive limited outreach, assuming zero effort outside the battleground provides a tractable approximation that closely matches the ob-

served spending distribution. Summary statistics are reported in Appendix C.3. The specific battleground states identified in each election year are shown in Figure 3 for 2008, 2012, 2016, and 2020, with the classification updated separately for each cycle.

Figure 3: Battleground States Across 2008 - 2020 Presidential Elections

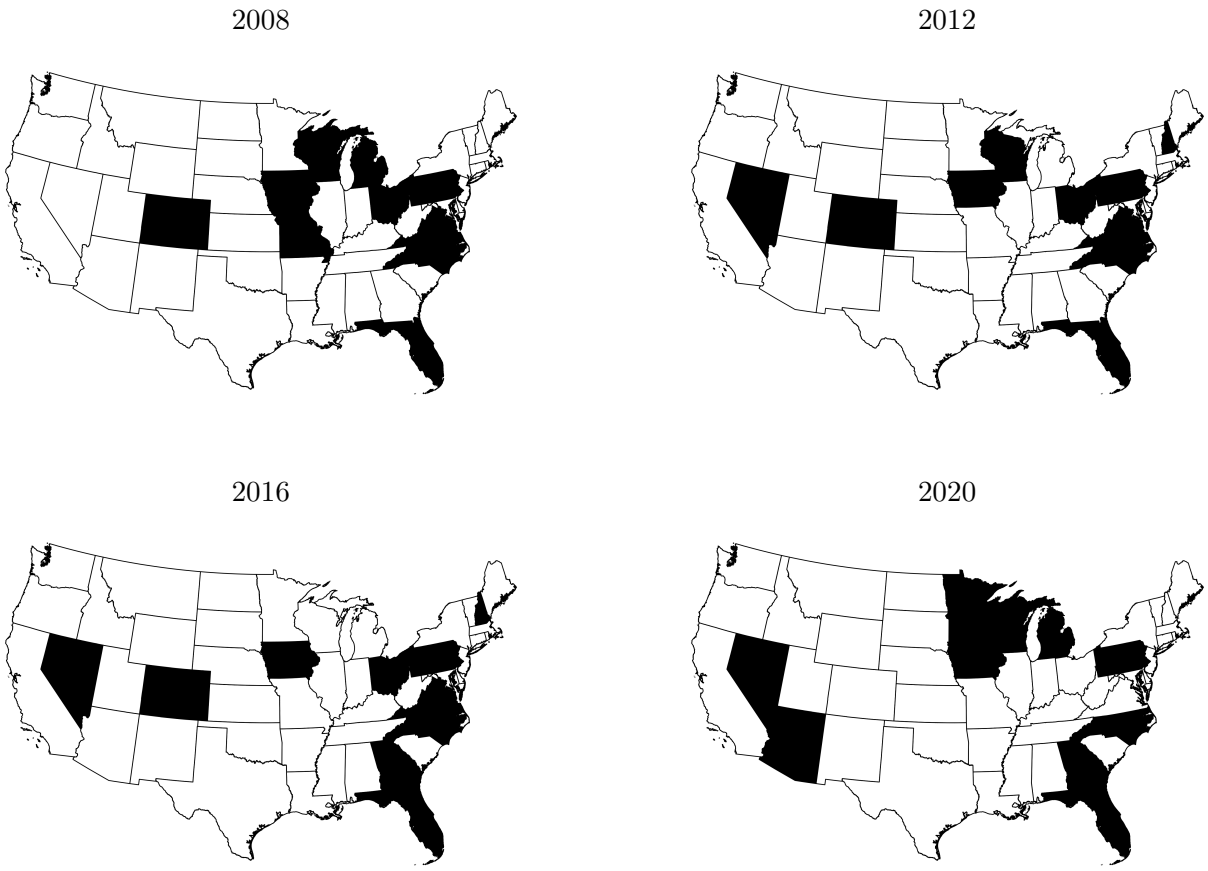

Notes: Each subfigure shows the battleground states (labeled in black) for the respective election year. The states are identified based on the total television advertising expenditures by both parties, with the top ten states by spending classified as battlegrounds.

While the Wesleyan Media Project provides detailed data on television advertising within battleground states, it does not capture the full scope of campaign mobilization. Television is a key channel, but campaigns also engage in digital outreach, direct mail, telemarketing, and travel, with meaningful variation across parties in how resources are allocated. Relying solely on television data would understate total campaign effort, and could bias comparisons

if one party disproportionately invests in non-television channels.

To address this, I use the Wesleyan data as a base to identify mobilization efforts within battleground states, and then scale these figures to recover total campaign budgets. This is done using Federal Election Commission (FEC) itemized disbursement records, which provide a comprehensive view of campaign spending by category.

Although the FEC includes purpose descriptions for each transaction, it does not report where the activity occurs. Vendor addresses are recorded, but reflect the firm's location, not the target of mobilization. For instance, an advertising agency based in California may deploy ads in Michigan.

To isolate direct voter contact, I classify each transaction into one of five mutually exclusive mobilization categories—media, digital, print, telemarketing, and travel— based on keyword matching in the purpose description field (see Appendix D). This procedure excludes overhead expenditures such as salaries, legal fees, and compliance costs.

Let FEC_{ptc} denote total spending by party p in category c during year t. I compute the national share of spending allocated to each category (denoted as ϕ_{ptc}) as:

$$\phi_{ptc} = \frac{\text{FEC}_{ptc}}{\sum_c \text{FEC}_{ptc}}.$$

Assuming that television ads correspond to the "media" category, I infer total effort in battleground state s as:

$$E_{spt} = \frac{\text{TV}_{spt}}{\phi_{pt\,\text{media}}},$$

and total party-level effort as:

$$E_{pt} = \sum_{s \in BG_t} E_{spt}.$$

Where BG_t is the set of battleground states for election year t. Table 4 reports the estimated composition of campaign disbursements, while Table 5 shows total inferred mobilization budgets by party and year, aggregated over battleground states.

To isolate general-election activity, I restrict the Wesleyan sample to ads aired on or after August 1 of each election year. By this point, major-party nominees are finalized, primaries are complete, and campaigns begin targeting the general electorate.

Table 4: Share of Disbursements by Category and Party

	2008		2012		2016		2020	
Category	Dem	Rep	Dem	Rep	Dem	Rep	Dem	Rep
Media	0.704	0.774	0.744	0.674	0.772	0.431	0.615	0.472
Online / digital	0.079	0.000	0.108	0.123	0.073	0.387	0.367	0.441
Print / mail	0.060	0.025	0.035	0.055	0.005	0.010	0.003	0.022
Phone / text	0.024	0.068	0.047	0.044	0.004	0.009	0.004	0.000
Travel / event	0.133	0.133	0.066	0.105	0.145	0.164	0.011	0.064

Notes: Entries report the fraction of total *itemized operating disbursements* made by the two major-party presidential candidate committees that falls in each spending category during the indicated election cycle. Shares are calculated from Federal Election Commission (FEC) operating-expenditure files after (i) restricting to general-election spending by the principal presidential committees, and (ii) assigning each transaction to a single category using keyword matching (see Appendix D for details). Columns sum to one within each party-year.

Table 5: Total Campaign Spending in Battleground States by Party (2020 USD)

Year	2008	2012	2016	2020
Democratic	257,000,000	349,000,000	153,000,000	392,000,000
Republican	158,000,000	221,000,000	155,000,000	416,000,000

Notes: Total estimated campaign expenditures in battleground states, aggregated at the party level. Values are in constant 2020 dollars and reflect the inferred mobilization budget constructed from television advertising data and scaled using national disbursement shares.

5 Results

Figure 4 and Table 6 summarize the estimated coefficients, showing how demographic and institutional covariates shape both baseline political alignment and the cost of voting. They also report the coefficients that map to key structural parameters: the utility gain from campaign effort (θ), the parameters of the perceived voting efficacy function (α_1, α_2), and the standard deviations of the county- and state-level shocks (σ_j, σ_s).[9]

5.1 Estimated Coefficients for Partisan Alignment

The μ coefficients govern counties' baseline political alignment ($\mu_{j_s t} = \beta_{\mu t} \cdot X^{\mu}_{j_s t}$) where $X^{\mu}_{j_s t}$ is a vector of county-level demographic and socio-economic variables. Negative values shift counties toward a Democratic baseline, while positive values shift them toward a Republican baseline. Figure 4 shows the estimated coefficients for the μ component, along with 95% confidence intervals. As can be seen, most coefficients are statistically different from zero. Further, the signs typically align with well-established empirical patterns. Counties with higher rates of college or advanced degree attainment, and larger Black and Hispanic population shares, lean more Democratic. In contrast, counties with higher proportions of white residents or lower educational attainment tend to lean more Republican. These results serve as a useful validity check on the model's estimates.

5.2 Estimated Coefficients for the Cost of Voting

Table 6 reports how institutional and socio-economic variables relate to the cost of voting. The cost function

$$c_{j_s t} = \exp(\beta_c \cdot X^{c}_{j_s t})$$

[9]For ease of interpretation, coefficients for the μ component are shown graphically in Figure 4, while the full set of estimates appears in Tables 16–19 in the Appendix.

maps these characteristics into a latent participation cost. Higher polling place congestion in both urban and rural areas is associated with significantly greater costs, while counties with higher median income or employment tend to face lower estimated costs.

Institutional rules show more mixed effects. States that require voters to register further in advance exhibit higher participation costs, consistent with evidence that easing registration requirements can raise turnout (Grumbach and Hill, 2022; Burden et al., 2014). In contrast, voter ID laws and polling hours have small and statistically insignificant coefficients. The null effect of ID laws aligns with recent findings that such requirements have limited impact on turnout (Cantoni and Pons, 2021).

Figure 4: Estimated Coefficients for the μ Component

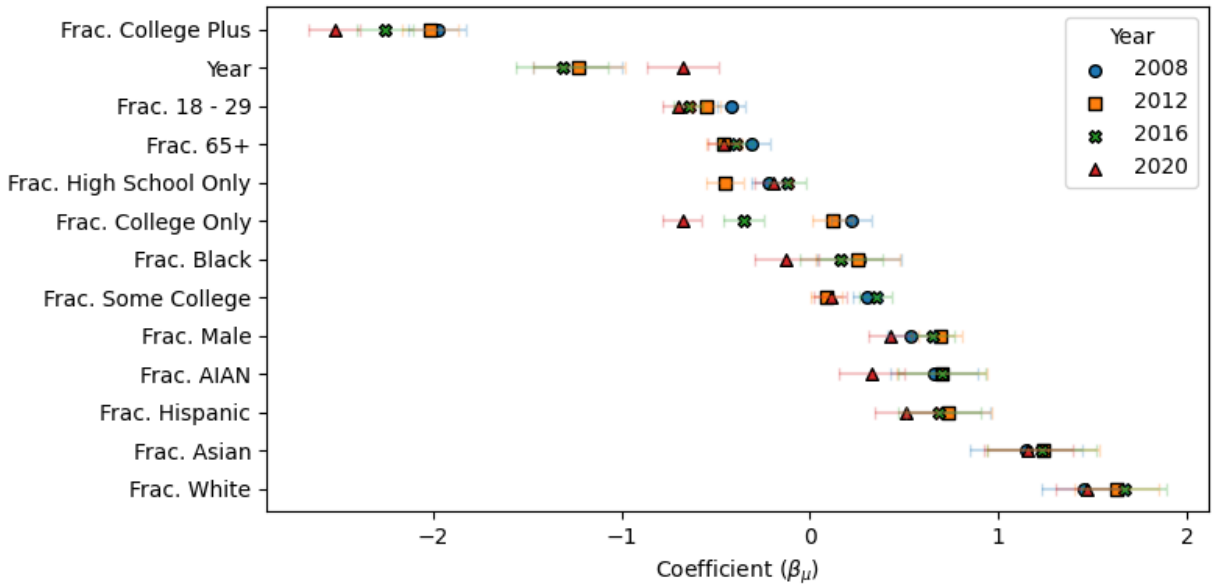

Notes: The plot shows the estimated coefficients for the model, with 95% confidence intervals. The coefficients represent the impact of each factor on the political alignment of counties, entered as $\mu_{j_s} = \beta_\mu \cdot X^{\mu}_{j_s}$. The x-axis shows β_μ coefficients, while the y-axis lists the variables. AIAN is the American Indian and Alaska Native population share.

Table 6: Estimated Coefficients

Parameter	Variable	Coefficient	Std. Error	P-Value
Cost of Voting	Constant	-0.481	0.020	0.000
	Voter ID laws	-0.011	0.007	0.107
	Urban polling density	0.457	0.026	0.000
	Rural polling density	0.361	0.024	0.000
	Employment rate	-0.678	0.027	0.000
	Polling hours	0.070	0.016	0.000
	Registration deadlines	0.198	0.006	0.000
Voting Efficacy	Constant 1	1.772	0.094	0.000
	Constant 2	-3.054	1.563	0.051
Campaign Effort	θ	-3.092	0.055	0.000
Std. Dev. County Shocks	σ_η	-1.991	0.004	0.000
Std. Dev. State Shocks	σ_δ	-1.525	0.068	0.000

Notes: All coefficients are from the maximum likelihood estimation. Cost-related covariates are expressed in standardized units (e.g., percentiles or logs).

5.3 County-Level Political Alignment

Figure 5 plots the estimated distribution of county-level baseline partisan preferences, $\mu_{j_s t}$, defined as the expected utility differential favoring the Republican candidate in county j of state s, prior to the realization of campaign effort or shocks.

In both battleground and non-battleground states, the distribution of $\mu_{j_s t}$ is unimodal and right-skewed, with average values of 0.10 and 0.11 respectively. These estimates are not population-weighted, so they reflect geographic rather than electoral mass and overrepresent rural, Republican-leaning counties.

Notably, the left tail is longer than the right in both cases. This implies that while most counties lean slightly Republican, the strongest partisan alignments tend to favor Democrats.

5.4 County-Level Cost of Voting

Figure 6 plots the estimated cost of voting across all U.S. counties, separately for battleground and non-battleground states. Table 7 summarizes the key statistics. The cost measure shown here, $c_{j_s t} = \exp(\beta_c \cdot X^c_{j_s t})$, reflects institutional and socioeconomic barri-

Figure 5: Estimated μ for all counties

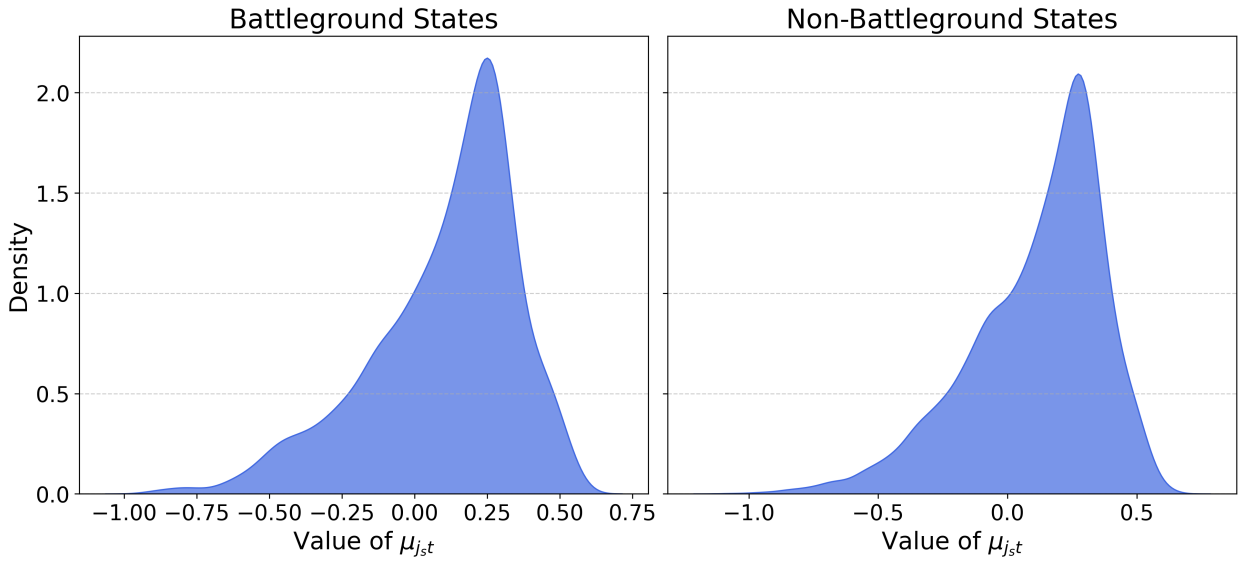

Notes: The figures show the estimated baseline county-level utility in favor of the Republican candidate (μ_{jst}). The left panel shows the estimated μ_{jst} for counties in battleground-states (where equilibrium effort levels are calculated), while the right panel shows the corresponding estimates for non-battleground states.

ers to participation, but does not yet account for perceived voting efficacy. I turn to that adjustment in the next section.

In the raw cost estimates, average values are quite similar across the two groups: the mean is 0.575 in battleground counties and 0.588 in non-battlegrounds, with similar dispersion (standard deviation = 0.073 vs. 0.064). The similarity reflects the fact that institutional voting costs, such as registration deadlines and polling congestion, do not systematically differ by state competitiveness.

However, this picture shifts once we account for perceived voting efficacy. Because efficacy is systematically lower in non-battleground states, the effective cost of voting is higher in those areas. The next subsection formalizes this adjustment and shows that, after incorporating efficacy, the gap in participation costs between battleground and non-battleground counties widens considerably.

Table 7: Summary Statistics for Estimated Cost of Voting by State Type

	Cost		Effective Cost	
Statistic	BG	Non-BG	BG	Non-BG
Mean	0.575	0.588	0.578	0.616
Std. Dev.	0.073	0.064	0.073	0.070
Min	0.359	0.369	0.361	0.376
Max	0.824	0.813	0.824	0.881

Notes: Raw cost is computed as $c_{j_s} = \exp(\beta_c \cdot X_{j_s}^{(c)})$. The effective cost adjusts for perceived voting efficacy, $c_{j_s}/p(\sigma_{s,D}, \sigma_{s,R})$, as defined in the turnout equations (Equations (4)–(5)). Lower values of p raise the effective cost, dampening turnout incentives. BG = Battleground states; Non-BG = Non-battleground states.

5.5 Perceived Voting Efficacy

As defined in Section 2, voting efficacy is modeled as a rescaled logistic function:

$$p(\kappa_s) = \frac{\tilde{p}(\kappa_s) - \tilde{p}(0)}{\tilde{p}(1) - \tilde{p}(0)}, \qquad \text{where} \quad \tilde{p}(\kappa_s) = \frac{1}{1 + \exp\left(-\alpha_1(\kappa_s - \alpha_2)\right)}$$

where κ_s measures the competitiveness of the state-level contest.

Maximum likelihood estimation yields

$$\widehat{\alpha}_1 = 5.88, \qquad \widehat{\alpha}_2 = 0.05,$$

implying a steep but continuous relationship between competitiveness and perceived voting efficacy that saturates quickly, as seen in Figure 7. Perceived voting efficacy reaches 0.90 by $\kappa_s = 0.51$ and exceeds 0.99 by $\kappa_s = 0.84$. This pattern suggests that voters perceive diminishing marginal returns to electoral closeness; once a contest is already moderately competitive, any further increase in closeness does little to change perceived voting efficacy.

Because efficacy scales the cost of voting as $c_{j_s}/p(\kappa_s)$, even small differences in κ_s materially affect turnout incentives. For instance, at a baseline cost $c_{j_s} = 0.59$ (the mean in non-battleground states), the ratio of cost to perceived voting efficacy is:

Figure 6: Estimated cost for all counties

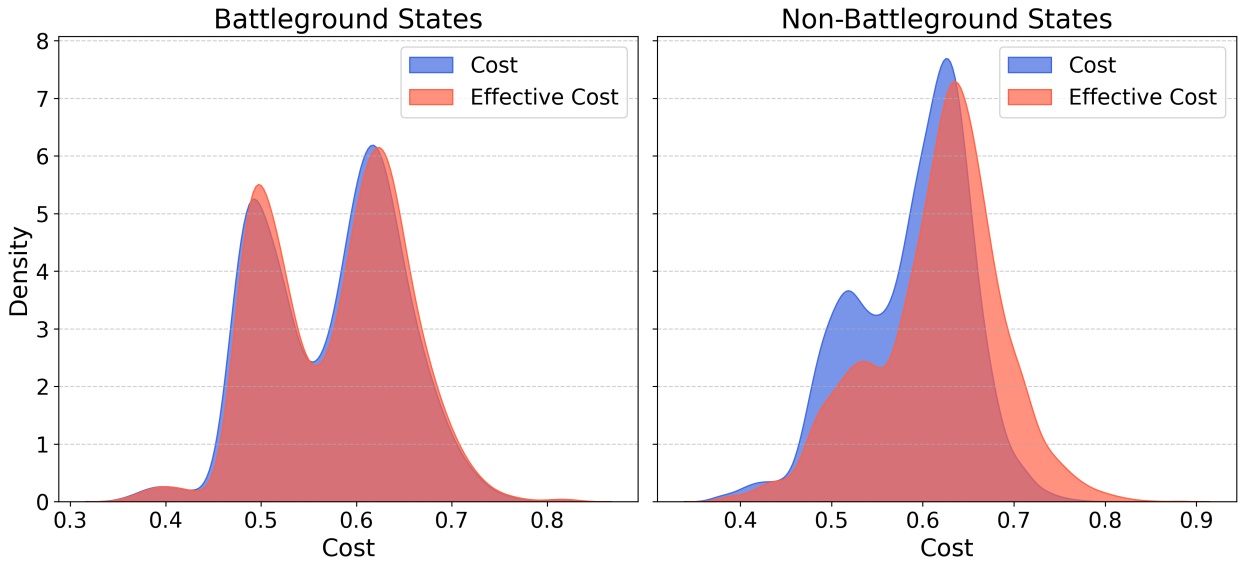

Notes: The figures show the estimated cost of voting for all counties. The left panel shows the estimates for counties where the equilibrium effort levels are calculated, while the right panel shows the corresponding estimates for non-competitive states.

$$\frac{0.59}{0.90} \approx 0.66 \qquad \text{vs.} \qquad \frac{0.59}{1.00} = 0.59,$$

implying a 12 percent increase in the effective cost of participation.

This adjustment is negligible in battleground states, where high competitiveness keeps perceived voting efficacy near its upper bound of 1.0. In non-battleground states, however, the effect is substantial. According to Table 7, accounting for lower efficacy raises the mean effective cost from 0.59 to 0.62. The impact is even more pronounced for the highest-cost voters, where the maximum effective cost rises from 0.81 to 0.88.

5.6 Equilibrium Campaign Effort

In equilibrium, candidates $q \in \{D, R\}$ choose how to allocate campaign effort across battleground states to maximize their probability of winning the national election. The impact of

Figure 7: Estimated Voting Efficacy Function

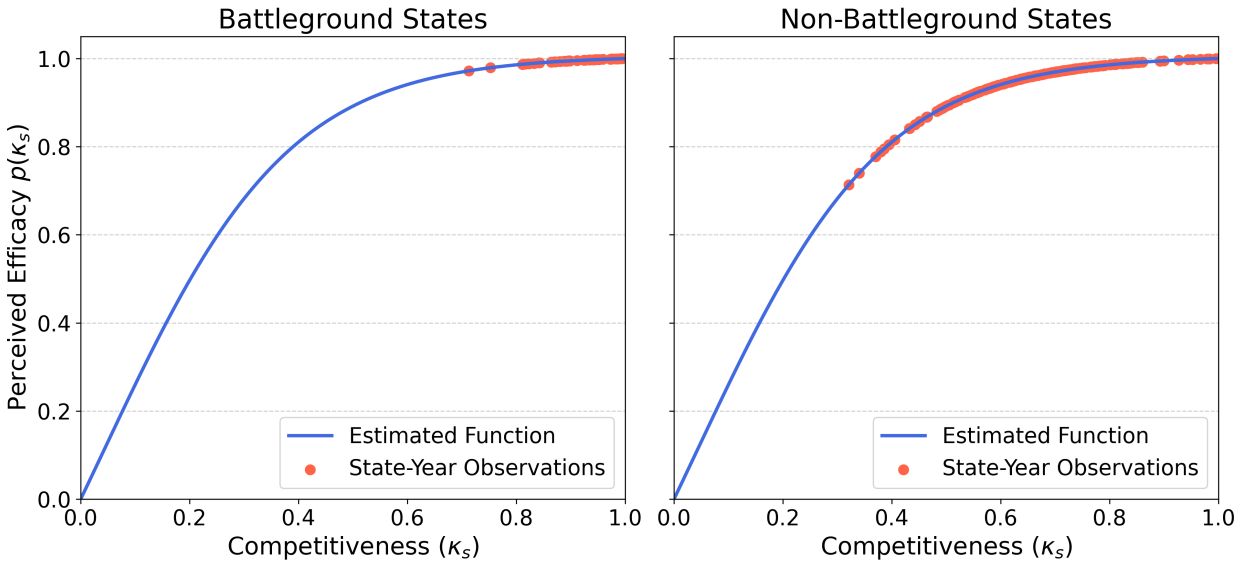

Notes: Estimated perceived voting efficacy as a function of the turnout ratio. Blue lines plot (2.1.3) evaluated at $(\widehat{\alpha}_1, \widehat{\alpha}_2) = (5.88, 0.05)$. Red circles are state-year observations.

effort on the utility differential is governed by the function

$$m(e_{s,q}, e_{s,-q}) = \theta\, e_{s,q}^{1/\gamma} - \psi\, \theta e_{s,-q}^{1/\gamma}.$$

A grid search over values of γ and ψ yields estimates $\hat{\gamma} = 2.0$ and $\hat{\psi} = 0.25$ as the best-fitting parameters. The corresponding maximum likelihood estimate for θ is $\hat{\theta} = 0.046$.

This results in an average utility gain of approximately 0.056 for Democratic partisans, and 0.047 for Republicans. To understand the magnitude of this effect, I compute the ratio of the utility gain from effort to the baseline partisan utility difference, $m(e_{s,q}, e_{s,-q})/|\mu_{j_s}|$, for each county. For the median county, this ratio is approximately 0.10 for both parties, indicating that for a typical county, campaign mobilization increases the effective partisan preference by a notable 10% relative to its baseline.

However, the distribution of this effect is highly skewed. The mean of the ratio is a much larger 0.67 for Democrats and 0.58 for Republicans. This large gap reveals that the impact of campaign effort is exceptionally strong in a subset of counties with very low initial partisan

leanings ($|\mu_{j_s}| \approx 0$). These politically marginal counties are therefore the most susceptible to being reshaped by campaign spending. Whether this translates into additional votes, however, depends on the other components of a voter's decision: the participation cost (c_{j_s}) and the competitiveness of the statewide contest ($p(\kappa_s)$), in addition to the county and state-level shocks.

To quantify marginal responsiveness to effort, I compute county-level elasticities of turnout with respect to campaign effort:

$$\varepsilon_{j_s,q,q} = \left(\frac{\partial \sigma_{j_s,q}}{\partial e_{s,q}}\right)\left(\frac{e_{s,q}}{\sigma_{j_s,q}}\right), \qquad \varepsilon_{j_s,q,-q} = \left(\frac{\partial \sigma_{j_s,q}}{\partial e_{s,-q}}\right)\left(\frac{e_{s,-q}}{\sigma_{j_s,q}}\right), \qquad q, -q \in \{D, R\}.$$

More details on the derivation can be found in Appendix F.1. Table 8 reports the average elasticity for both same and cross-elasticities across all battleground states and election years, weighted by the county population.

Table 8: Mean elasticity of turnout with respect to own and opponent effort

	Effort (D)	Effort (R)
Turnout (D)	0.053	−0.012
Turnout (R)	−0.011	0.040

A 10% increase in Democratic effort raises Democratic turnout by roughly 0.53% and lowers Republican turnout by 0.11%. Conversely, a 10% increase in Republican effort raises Republican turnout by about 0.40% and lowers Democratic turnout by 0.12%. The finding that campaigns primarily affect their own voters, with smaller impacts on their opponents, suggests the main channel of influence is mobilization rather than persuasion. Further, these modest elasticities suggest at current levels of campaign spending, the marginal voter is relatively unresponsive to additional effort. This finding, however, stands in contrast to the meaningful influence campaign effort exerts across the entire spectrum of voters, as will be discussed in Section 7

6 Model Fit and Validation

The credibility of the structural model depends not just on internal coherence, but on its ability to reproduce empirical patterns that were not mechanically imposed during estimation. This section validates the model along two critical dimensions. First, on the voter's side, I show the model's core behavioral mechanism, the turnout response to intrinsic competitiveness, is quantitatively consistent with reduced-form evidence from a border discontinuity design. Second, having established this behavioral foundation, I demonstrate the model's predicted effort allocations from the candidate side of the model align closely with observed candidate activity.

6.1 Matching Turnout Responses to Competitiveness

The structural model predicts intrinsic competitiveness has a direct effect on turnout, independent of campaign effort. This effect operates through the estimated voting efficacy function: as competitiveness increases, a voter's perceived efficacy rises, making turnout more attractive; when competitiveness declines, perceived efficacy falls, discouraging participation. To evaluate whether the model generates a quantitatively accurate response to competitiveness, I compare its predicted turnout effect to reduced-form estimates that isolate comparable variation.

To isolate the effect of competitiveness on turnout within the model, I simulate a counterfactual in which a county is hypothetically placed into a state where the election is expected to be exactly tied ($\kappa = 1$). The county's own characteristics (partisanship, voting costs, campaign effort, and structural shocks) are held fixed, meaning only the perceived competitiveness of the statewide environment is altered. Solving equations 4 and 5 under this scenario for each non-battleground county yields an average turnout increase of 1.74 percentage points across non-battleground counties[10].

I assess the credibility of this prediction using a border discontinuity design that compares

[10] I do not use battleground states for this counterfactual as $p(\kappa_s) \approx 1$ for all of these states

turnout across adjacent counties in different states but within the same media market. These counties are exposed to identical political advertising and share similar demographics due to geographic proximity, but differ in statewide competitiveness. Because campaign advertising is held constant within a media market, the design offers a clean quasi-experimental test of how turnout responds to competitiveness.

A key advantage of the border discontinuity design is that it compares counties within the same media market, thereby holding constant exposure to broadcast advertising. Nonetheless, two potential threats to identification warrant careful consideration. The first is that ground operations such as field offices or campaign events may vary discontinuously across state lines. The second, is that digital advertising, unlike broadcast television, can be targeted at the state level, potentially violating the assumption of equal campaign exposure within media markets. Two robustness checks, discussed below, verify that these concerns do not drive the main results.

To create a direct counterpart to the model's prediction, I use the reduced-form estimates to calculate the implied turnout gain from moving a typical non-battleground state to a perfect tie. Using the mean non-battleground competitiveness ($\kappa \approx 0.694$) as the baseline, the estimates imply a turnout increase of 1.67 percentage points. This range is remarkably close to the model's predicted average effect of 1.74 percentage points, providing strong external validation for the model's competitiveness channel.

6.1.1 Border Discontinuity Design

Figure 8 illustrates the logic of the border discontinuity design using the Ohio–Kentucky media market. Counties on both sides of the state line receive the same political advertising but face different levels of statewide competitiveness. Identification relies on the assumption that, conditional on shared media exposure, turnout differences across the state border capture differences in electoral competitiveness.

The analysis uses the same county-level dataset as the structural estimation, restricted to

Figure 8: County-Level Variation in Competitiveness and Campaign Effort

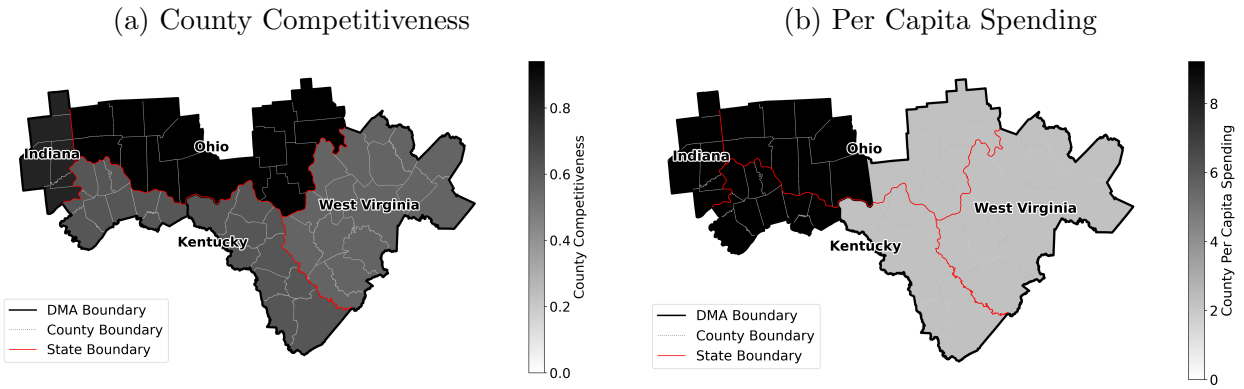

Notes: Plots show county-level variation in competitiveness and campaign effort within the Cincinnati and Charleston-Huntington DMAs in 2012. Both regions span multiple states (Indiana, Ohio, West Virginia, and Kentucky), allowing for comparisons across counties with shared media exposure but differing statewide electoral contexts. Panel (a) plots state-level competitiveness, assigned to each county based on its state-year. Panel (b) shows total per-capita campaign spending across counties, aggregated from candidate disbursements and media market ad data. As can be seen, competitiveness sharply varies across state borders, while per-capita spending is constant throughout the DMA. Each map overlays DMA boundaries (black), state borders (red), and county lines (gray).

counties located on state borders within media markets that span multiple states. Following Spenkuch and Toniatti (2018), counties bordering multiple neighbors are included once for each pairwise comparison. Counties that are split across multiple media markets are dropped. Only 8 counties are dropped for this reason, out of over 3,000 observations.

Balance tests (Appendix G.1) regress county-level covariates on an indicator for being in the higher-turnout state within each border pair, controlling for border-pair-by-year fixed effects. With the exception of a modest imbalance in Hispanic population share, none of the demographic, educational, or economic covariates show significant differences, supporting the validity of the identification strategy.

The main regression is conducted at the county-border pair-year level, with each observation representing a county matched to a neighboring county across a state border, and within the same media market. Turnout is regressed on state-level competitiveness using:

$$\text{Turnout}_{cpt} \;=\; \beta\,\kappa_{s(c)t} \;+\; \Gamma \cdot X_{cpt} \;+\; \delta_{pt} \;+\; \chi_{s(c)} \;+\; \varepsilon_{cpt} \tag{18}$$

where Turnout_{cpt} is the share of the voting-age population in county c, border pair p, and year t that cast ballots for either the Democratic or Republican presidential candidate. The key regressor, $\kappa_{s(c)t}$, is the state-level competitiveness of the state containing county c in year t, defined identically to the model as the ratio of trailing to leading party's two-party vote share.

The fixed effects δ_{pt} absorb all time-varying factors shared by counties within each border pair. Most importantly, this includes effects from campaign exposure, since paired counties lie within the same media market. State fixed effects $\chi_{s(c)}$ account for time-invariant state characteristics that may influence turnout. The vector X_{cpt} includes county-level controls for demographics, education, and economic conditions, which are included in robustness checks to ensure that the results are not driven by imbalances in these covariates. The coefficient β is therefore identified from cross-state variation in competitiveness $\kappa_{s(c)t}$ within each border pair, after conditioning on fixed effects that absorb campaign exposure and time-invariant state factors.

6.1.2 Results

Table 9 reports the main regression-discontinuity estimates. Column (1) presents the baseline specification without covariates. Column (2) adds Hispanic population share, the only variable flagged as imbalanced in the balance tests. Column (3) includes the full set of demographic and economic controls. Standard errors (clustered by county-pair) appear in parentheses. Across specifications, the estimated coefficients range from 0.054 to 0.064, and are statistically significant at the $p < 0.001$ level.

A series of robustness checks confirm the stability of these results. First, to address concerns that ground operations may vary across borders, I re-estimate the model excluding all counties with evidence of field offices or campaign events, identified using Federal Election

Commission data. The resulting estimates change only marginally, indicating that cross-border differences in field operations do not drive the main results. More details on selecting these counties and the results are provided in Appendix G.2.

A second consideration is that digital advertising, unlike broadcast television, can be targeted at the state level, potentially violating the assumption of equal campaign exposure within media markets. To account for this, I incorporate measures of per-capita digital spending from Facebook in 2020. The coefficient on competitiveness remains statistically significant, though larger in magnitude, likely reflecting the absence of state fixed effects in this specification. As a complementary test that circumvents this data limitation, I restrict the sample to the 2008 and 2012 elections, a period when digital advertising was significantly less prevalent, as shown in Table 4. In this early digital-era sample, the estimated effect of competitiveness is stable and remains highly consistent with the main results, ranging from 0.054 to 0.064. Taken together, these checks provide confidence that the findings are not an artifact of confounding from digital campaign spending.

Finally, I test whether results are sensitive to the measurement of competitiveness. Instead of realized vote shares, Appendix G.4 uses pre-election polling to construct κ_{st}. Although the estimated coefficients are smaller (0.029–0.030), they remain positive and highly significant, reinforcing that the turnout response is not an artifact of post-treatment measures of closeness.

These checks demonstrate that the estimated effect of competitiveness on turnout is both robust and substantively meaningful. The coefficients remain positive and highly significant across a range of specifications, samples, and alternative measures of competitiveness, providing confidence that the reduced-form relationship is genuine rather than an artifact of campaign outreach or measurement choice. This stability establishes a credible benchmark against which to evaluate the model's predicted turnout response.

Table 9: Effect of Competitiveness on Turnout

		(1)	(2)	(3)
Competitiveness	Estimate	0.054***	0.059***	0.064***
	(Std. Error)	(0.007)	(0.007)	(0.007)
Controls		No	Balance-Test Sig.	All
Border Pair by Year FE		Yes	Yes	Yes
State FE		Yes	Yes	Yes
Observations		8,126	8,126	8,126
R^2_{within}		0.003	0.156	0.534

Notes: Each column reports regression estimates of the effect of state-level competitiveness on turnout, measured at the county-border pair-year level. All models include border-pair-by-year fixed effects and state fixed effects. Columns (2) and (3) sequentially add controls for covariates flagged as imbalanced in the balance tests and the full set of demographic and economic covariates. Standard errors (in parentheses) are clustered at the county-pair level. Significance levels: * p<0.05, ** p<0.01, *** p<0.001.

Table 10: Estimated Turnout Effect of Moving from Average Non-Battleground Competitiveness to Full Competitiveness

	No Controls	Balance Sig. Controls	All Controls
Effect (pp)	1.67	1.82	1.98

Notes: Each entry reports the estimated increase in turnout (in percentage points) associated with raising competitiveness from the non-battleground state mean ($\kappa_{st} = 0.694$) to full competitiveness ($\kappa_{st} = 1$), based on the coefficients in Table 9.

6.1.3 Discussion

To compare with the model's prediction, I compute the implied turnout effect of increasing competitiveness from the non-battleground mean ($\kappa_{st} = 0.694$) to a perfect tie ($\kappa_{st} = 1$), a change of 0.306 units. Multiplying this difference by the estimated coefficient on competitiveness yields a predicted turnout increase of 1.67 percentage points using the baseline specification (see Table 10). This closely aligns with the model's predicted effect of 1.74 percentage points. The estimates remain stable across specifications, increasing to 1.82 and 1.98 percentage points when controlling for imbalanced covariates and then using the full set of controls, respectively. While the estimates are slightly larger when adding covariates, they still remain close to the model's prediction. The correspondence between model-implied and quasi-experimental estimates provides strong external validation of the recovered voting

efficacy function. It suggests that the model accurately captures the behavioral response to competitiveness and supports the credibility of the model's counterfactual predictions.

The magnitude of the reduced-form estimates is notable relative to existing causal evidence. In redistricting-based studies of U.S. legislative elections, increased competitiveness typically raises turnout by less than one percentage point (Ainsworth, Munoz, and Gomez, 2022; Fraga, Moskowitz, and Schneer, 2022). Field experiments that manipulate perceptions of closeness often find no detectable effect on turnout, casting doubt on pivotality-based mechanisms (Enos and Fowler, 2014). By contrast, Bursztyn et al. (2024) document sizable turnout effects in response to close polls in Swiss referenda, though in a different institutional context.

The present findings differ in two key respects. First, prior studies often examine lower-salience elections, such as local or state legislative contests, where electoral competitiveness may be less salient to voters. In contrast, this paper focuses on U.S. presidential elections, where political stakes are high and baseline engagement is substantially greater.

Second, many existing designs confound competitiveness with campaign activity. Redistricting studies, for example, assign voters to more competitive districts, but such districts are often the target of heightened mobilization efforts. In that setting, observed increases in turnout may reflect campaign outreach rather than competitiveness per se. The border discontinuity design avoids this conflation by comparing adjacent counties within the same media market, holding campaign exposure constant while allowing competitiveness to vary at the state level.

This setting allows for a clean test of the model's core behavioral prediction: that competitiveness, independent of campaign mobilization, can meaningfully affect turnout. The magnitude and robustness of the estimated response support this interpretation, demonstrating that competitiveness is not only salient but behaviorally consequential, even net of campaign intensity.

To ensure the results of the border discontinuity are not driven by idiosyncrasies of

the data, I conduct a complementary robustness check replicating the analysis of Spenkuch and Toniatti (2018). Their approach compares counties on opposite sides of a media market boundary but within the same state, thereby holding competitiveness constant while allowing television advertising exposure to vary. I recover similar estimates: campaign advertising has no discernible effect on turnout, while differences in partisan spending meaningfully affect vote shares. The close replication of their findings indicates that the turnout effects in my main design are not artifacts of sample selection or measurement, but reflect genuine effects of electoral competitiveness. See Appendix G.5 for details.

6.2 Matching Effort Allocations to Campaign Activity

Validating the candidate side of the model requires comparing its predicted equilibrium effort allocations to observed campaign behavior. A primary challenge in this validation lies in mapping the model's theoretical concept of effort to its empirical counterparts. The model, for reasons of parsimony and tractability, conceptualizes effort ($e_{s,p}$) as a single, unified resource. In practice, this corresponds to a diverse portfolio of campaign activities, including television, digital outreach, in-person visits, and ground operations like canvassing. The optimal mix of this portfolio likely varies by state as campaigns adapt to local conditions and the tools within this portfolio have different characteristics. Some, like ad buys, are highly divisible, while others, like candidate visits, are inherently lumpy.

The validation is further complicated by the fact that only a subset of these activities are captured in publicly available data: television, digital, and presidential visits. Consequently, the analysis necessarily compares the model's prediction of total effort against an incomplete, though varied, set of its real-world components.

Nevertheless, one should still expect the model's predicted effort allocations to align with various proxies for campaign effort. Where the model predicts high amounts of effort, one should observe substantial campaign activity. Further, given the model's assumption that effort is a single, divisible resource, it is expected that the model's predictions will align

more closely with proxies that reflect scalable and divisible campaign activities, such as digital and television advertising, while showing weaker correspondence with proxies that capture indivisible or "lumpy" activities, like in-person candidate visits.

With this in mind, Table 11 summarizes the fit of the model's predicted effort shares against three key proxies for campaign effort: television advertising, presidential candidate visits, and digital advertising. Digital advertising, sourced from OpenSecrets, is only available for the 2020 election cycle.

As anticipated by the framework, the model's predicted allocations align strongly with observed spending on large-scale media. The fit for television advertising is robust across all cycles, with correlations of $\rho = 0.802$ for Democrats and $\rho = 0.760$ for Republicans. The correspondence is even stronger for digital advertising, the most granular proxy, where the correlation between predicted state-level effort shares and observed spending in 2020 is $\rho = 0.907$ for Democrats and $\rho = 0.911$ for Republicans. The correspondence is weaker for the indivisible and logistically complex proxy of in-person presidential visits. For these events, the correlations fall to $\rho = 0.600$ for Democrats and $\rho = 0.539$ for Republicans.

Table 11: Fit statistics by effort proxy and party

Effort proxy	Correlation ρ		RMSE	
	Dem	Rep	Dem	Rep
Television ads	0.802	0.760	.044	.048
Presidential candidate visits	0.600	0.539	.070	.065
Digital ads (2020 only)	0.907	0.911	.053	.067

Correlation and RMSE values reflect the match between predicted effort shares and observed proxies by state. RMSE is computed as $\sqrt{\frac{1}{S}\sum_s(\hat{e}_s - e_s)^2}$ and reported in percentage points of total national share.

Figure 9 visualizes the match by year. Within each panel, states are ordered by predicted effort. Bars show the model's predicted effort alongside observed campaign activity from television advertising, in-person presidential visits, and, for 2020, digital advertising. High-priority states such as Florida consistently appear at the top, while low-priority states such

as Iowa remain at the bottom. This alignment provides strong validation for the model's predictions.

7 Discussion and Policy Implications

This section examines the broader implications of the model for voter turnout and campaign behavior. I begin by decomposing the turnout gap between battleground and non-battleground states into the contributions of electoral competitiveness, captured through perceived voting efficacy, and direct campaign effort. I then address the puzzle of why campaigns invest heavily in mobilization despite modest treatment effects found in field experiments, showing how the model reconciles low marginal returns with large aggregate effects on turnout. Finally, I use the model to evaluate a counterfactual public financing reform that equalizes and reduces campaign budgets, illustrating how institutional changes to candidate resources can reshape both participation and electoral competitiveness.

7.1 Disentangling Intrinsic Competition and Strategic Effort

On average, 62.6% of eligible voters in battleground states turn out in presidential elections, compared to just 56.5% in non-battleground states (see Table 3). According to the structural model, this 6.1 percentage point gap can be fully accounted for by differences in electoral competitiveness and campaign effort. To quantify their individual contributions, I simulate a sequence of counterfactuals, shutting down each mechanism in turn while holding all other components of the model fixed. The results show that roughly one-third of the gap is attributable to differences in intrinsic competitiveness, while the remaining two-thirds are due to strategic campaign effort.

Figure 9: Estimated Effort Levels by Election Year

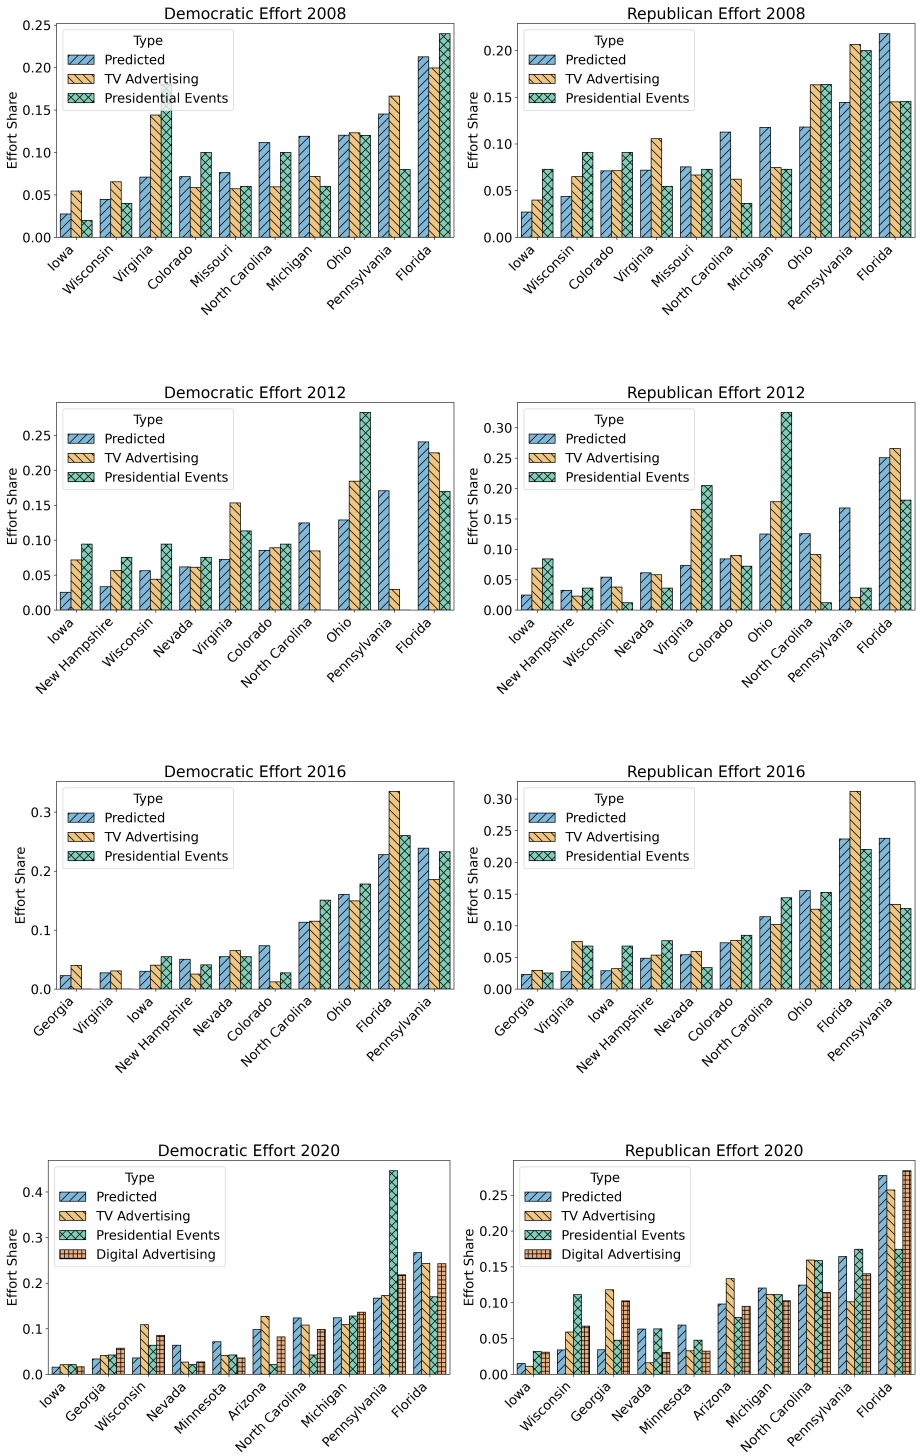

Notes: Each panel shows the model's predicted effort shares alongside observed campaign activity from television advertising, in-person presidential visits, and digital advertising (for 2020). Effort types are distinguished by both color and hatch patterns (see legend). States are ordered by predicted effort, with the total national share summing to 100%.

7.1.1 Battleground States

I begin from the observed equilibrium in battleground states and conduct a two-step counterfactual. First, I eliminate campaign effort by setting $e_{sD} = e_{sR} = 0$ in the turnout equations (4) and (5), while keeping each state's equilibrium level of voting efficacy $p(\kappa_s)$ fixed. This shuts down the direct effect of effort but retains high levels of competitiveness, characteristic of battleground states. This adjustment reduces turnout from 62.6% to 58.0%, a decline of 4.6 percentage points. In absolute terms, turnout falls from 40.9 million votes to 38.0 million on average, a loss of 2.9 million votes per election cycle.

Second, continuing to hold effort at zero, I reduce voting efficacy by replacing $p(\kappa_s)$ with the average level observed in non-battleground states, 0.94. This further lowers turnout to 56.1%, a decline of 1.9 percentage points relative to the previous step. In absolute terms, turnout falls on average to 36.6 million votes, a loss of an additional 1.4 million.

7.1.2 Non-Battleground States

I apply a reverse counterfactual to evaluate how turnout in non-battleground states would respond if they were assigned the characteristics of battleground states. These states begin from an average equilibrium turnout of 56.6%, reflecting both low voting efficacy and no campaign effort (see Table 12).

First, I increase voting efficacy by setting $p(\kappa_s) = 0.99$, the average level in battleground states, while holding campaign effort at zero. This raises turnout to 58.6%, a gain of 2.0 percentage points. Introducing campaign effort at the average battleground level yields a further 3.6-point increase, bringing turnout to 62.3%.

The same pattern holds in absolute terms: increasing voting efficacy alone adds 2.8 million votes, and introducing campaign effort adds another 6.1 million, for a total increase of nearly 9 million votes across non-battleground states. Roughly one-third of the increase is attributable to competitiveness, and two-thirds to mobilization.

7.1.3 Summary

The model accounts for the full turnout gap by adjusting only two components: electoral competitiveness and campaign effort. When applied to non-battleground states, raising both mechanisms to the levels observed in battleground states, the counterfactual predicts turnout of 62.3 percent, which is nearly identical to the 62.6 percent observed in battleground states. When applied in reverse, reducing competitiveness and effort in battleground states to non-battleground levels, the model predicts turnout of 56.1 percent, closely matching the 56.6 percent observed in non-battleground states. This symmetry suggests that these two forces are sufficient to explain observed differences in turnout across these two groups. Further, both mechanisms contribute significantly to the turnout gap. While the effect of intrinsic competitiveness is smaller than that of strategic effort, its impact is still substantial, accounting for roughly one-third of the overall difference, or 2.0 percentage points of turnout. Ignoring either mechanism would substantially mischaracterize the drivers of turnout, highlighting the importance of including both mechanisms in any comprehensive electoral model.

Table 12: Counterfactual Decomposition of Turnout Gap

Scenario	Turnout	Δ from Eqm	Δ from Baseline	Share of Total	Votes (M)
Battleground States					
Equilibrium	0.626	—	—	—	40.90
No Effort, same p_s	0.580	-0.046	-0.046	0.699	37.99
No Effort, Low p_s	0.561	-0.019	-0.065	0.301	36.63
Non-Battleground States					
Baseline (Low p_s, No Effort)	0.566	—	—	—	92.78
High p_s, No Effort	0.586	+0.020	+0.020	0.356	95.53
High p_s, With Effort	0.623	+0.036	+0.056	0.644	101.65

Notes: Turnout is the share of eligible voters. Votes are in millions. "Δ from Eqm" refers to the change from the previous row in each panel. "Δ from Baseline" refers to the change from the first row in each panel. "Share of Total" refers to the fraction of the total turnout change explained by each component.

7.2 The Cost of Mobilization

Field experiments consistently find modest marginal effects of campaign outreach on turnout (e.g., Aggarwal et al., 2023; Spenkuch and Toniatti, 2018; Gerber et al., 2011). Yet presidential campaigns routinely invest hundreds of millions of dollars in battleground mobilization. This apparent disconnect raises the possibility that campaigns could reduce outreach without jeopardizing their electoral prospects.

However, most field experiments are conducted in environments where outreach is already widespread, so they recover only local treatment effects near saturation. The model captures this: at estimated equilibrium levels, the marginal cost of mobilizing an additional vote exceeds $250. But this high marginal cost does not imply that campaigns can safely reduce effort. The overall effect of mobilization remains quite large. Cutting back would depress turnout and entail serious electoral consequences.

To quantify this, let $V_{s,p}(e_{s,D}, e_{s,R})$ denote the model-implied vote total for party p in state s, given campaign efforts $e_{s,D}$ and $e_{s,R}$. The total effect of mobilization is given by:

$$\Delta V^*_{s,p} = V_{s,p}(e^*_{s,D}, e^*_{s,R}) - V_{s,p}(0, 0),$$

where $e^*_{s,p}$ is the estimated equilibrium effort level. Aggregating across all battleground states yields

$$\Delta V^*_p = \sum_{s \in BG} \Delta V^*_{s,p}.$$

Table 13 reports the additional votes generated by mobilization in each election from 2008 to 2020, along with corresponding turnout increases. On average, mobilization added 2.9 million votes, or about 4.4 percentage points, to turnout in battleground states.

Table 13: Votes Added and Turnout Increase from Campaign Mobilization

Year	Votes Added (Millions)			Turnout Increase (Percentage Points)		
	Dem	Rep	Total	Dem	Rep	Total
2008	1.66	1.04	2.70	2.46	1.54	4.00
2012	1.75	1.16	2.91	2.88	1.91	4.79
2016	1.06	1.09	2.15	1.59	1.64	3.23
2020	1.88	1.96	3.84	2.69	2.81	5.50
Avg	1.59	1.31	2.90	2.41	1.98	4.38

Notes: "Votes Added" reports the estimated number of additional votes due to campaign mobilization, relative to a zero-effort counterfactual. "Turnout Increase" is expressed in percentage points relative to the voting-age population in battleground states.

These vote gains translate into meaningful shifts in Electoral College outcomes. In a counterfactual where the Democratic Party eliminated its mobilization effort while Republicans maintained theirs, Democrats would have lost nearly 1.5 million votes per cycle in battleground states. As a result, they would have lost three of the four elections between 2008 and 2020, rather than winning three. Even their 2008 victory would have narrowed sharply, with just 302 electoral votes instead of 365.[11]

To assess what field experiments capture, I calculate the marginal cost of generating an additional vote at equilibrium effort levels. This metric corresponds to the local treatment effect of a small change in campaign intensity, precisely what a randomized field experiment would estimate. The ACV for party p in state s, is given by

$$ACV_{s,p} = \frac{e^*_{s,p}}{\Delta V^*_{s,p}},$$

[11]Under the counterfactual, Democratic electoral vote totals fall to 302 (2008), 227 (2012), 223 (2016), and 227 (2020), compared to actual outcomes of 365, 332, 232, and 306.

while the marginal cost per vote (MCV) is:

$$MCV_{s,p} = \left(\frac{\partial V_{s,p}}{\partial e_{s,p}} \right)^{-1}.$$

Details on the MCV calculation are in Appendix H.

To analyze how costs vary with campaign scale, I simulate counterfactuals that scale each party's budget proportionally. Figure 10 plots the results. The left panel shows total votes added as a function of effort. The right panel reports average and marginal costs per vote, averaged across battleground states and years.

Figure 10: Average and Marginal Costs per Vote Across Effort Scales

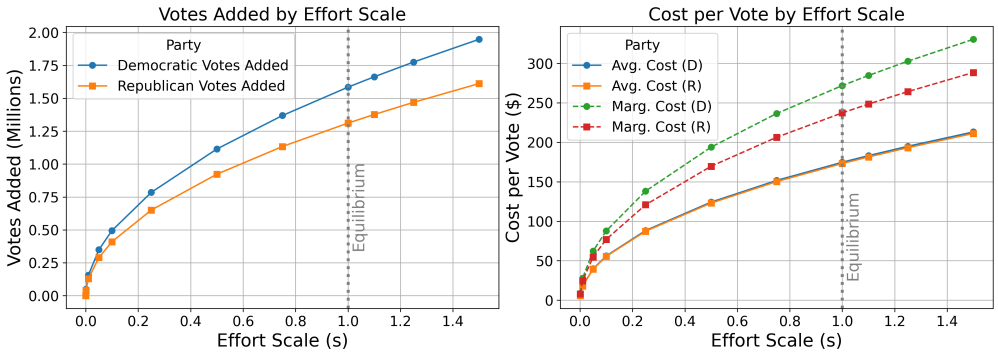

Notes: Left panel: total votes added from campaign mobilization as a function of scaled effort. Right panel: average and marginal cost per vote.

At low effort levels, mobilization is highly cost-effective, with average costs below $100 per vote. But costs rise steeply with scale: at equilibrium, marginal costs exceed $200 and average costs surpass $150. These diminishing returns explain the small effects seen in field experiments, even as aggregate effects remain large.

Taken together, these results reconcile two stylized facts. Marginal effects of mobilization are small at observed effort levels, consistent with field experiment estimates. Yet the aggregate impact remains substantial, warranting continued investment. Eliminating campaign effort would induce large turnout declines and considerable electoral losses.

7.3 Counterfactual: Public Financing Reform

To demonstrate the model's utility, I simulate a counterfactual public financing reform that caps each party's campaign budget at the inflation-adjusted equivalent of the 2008 public grant of $84.1 million. Under federal law, major party nominees have been eligible for such grants if they accept spending limits and forgo private fundraising. In 2008, the Republican nominee accepted the grant while the Democratic nominee declined it, and no major party candidate has participated since, relying instead on privately raised funds that now total several hundred million dollars. Note this simulates a more extreme version of public financing, assuming that total campaign spending is capped at the public grant level.

The reform affects campaigns through two channels. First, it equalizes budgets, a change especially consequential in 2008 and 2012 when Democrats held a significant spending advantage. Second, it reduces the overall level of resources, cutting spending by more than $200 million per candidate in some cycles. For each election cycle, I solve for the optimal allocation of campaign effort under the new budget constraints and compare electoral outcomes, turnout, and competitiveness to the observed baseline.

Table 14 shows that during the years when Democrats had a large spending edge, the reform closes the electoral college gap by 15-30 votes. In 2016 and 2020, when spending was more balanced, the reform has no effect on electoral outcomes.

Table 14: Electoral College Votes: Observed vs. Reform

	2008	**2012**	**2016**	**2020**
Democratic EVs (Observed)	364	332	233	306
Democratic EVs (Reform)	349	303	233	306
Change in Dem. EVs	*-15*	*-29*	*0*	*0*
Republican EVs (Observed)	174	206	305	232
Republican EVs (Reform)	189	235	305	232
Change in Rep. EVs	*+15*	*+29*	*0*	*0*

Notes: Observed values report actual Electoral College outcomes. Reform values are simulated outcomes under a public financing system that caps each party's campaign budget at the inflation adjusted equivalent of the 2008 public grant of $84.1 million. Positive changes for one party correspond to equal and opposite changes for the other.

Table 15 clarifies the mechanism behind these results. In 2008 and 2012 Democrats consistently outspent Republicans across battleground states, whereas by 2016 and 2020 both parties devoted roughly equal resources to campaigning. Under the reform, the Democratic share of total spending drops to 50 percent in all four cycles, as candidates match each other's effort under the new budget constraints. As a result, the reform eliminates Democrats' earlier resource advantage, reducing their electoral votes in 2008 and 2012, but has little effect in 2016 and 2020 when spending was already more balanced.

The reform also narrows margins while lowering participation. Average victory margins decline by 1.2 percentage points in 2008 and 1.4 points in 2012, yet this comes at the cost of reduced turnout. As shown in Table 15, turnout falls by an average of 1.08 million votes per election, with the largest drop of nearly 2 million in 2020.

Table 15: Effects of Public Financing Reform in Battleground States

	2008	**2012**	**2016**	**2020**
Dem. Effort Share (Observed)	62%	61%	50%	49%
Dem. Effort Share (Reform)	50%	50%	50%	50%
Change in Total Votes (millions)	-0.79	-1.17	-0.41	-1.94
% States More Competitive	80%	90%	60%	30%
Reduction in Victory Margin (pp)	1.2	1.4	0.04	0.2

Notes: Table reports simulated effects of a public financing reform that caps each party's campaign budget at the inflation adjusted equivalent of the 2008 public grant of $84.1 million. Effort shares reflect the average Democratic share of total campaign spending across battleground states (summing to 100% with the Republican share). Changes in votes are reported in millions. Competitiveness is measured as the share of battleground states with smaller victory margins under the reform. Reductions in victory margins are average state-level declines, expressed in percentage points.

Because the model explicitly accounts for abstention, it can trace how budget constraints alter both the intensity of mobilization and the decision to participate at all. The reform's effects reveal a central tradeoff: it heightens competitiveness by eliminating spending disparities, but it also reduces overall turnout as campaigns scale back their efforts. By modeling both strategic decision making by voters and campaigns, the model lets us capture this nuanced dynamic.

8 Conclusion

This paper develops a novel structural model of turnout in U.S. presidential elections that embeds voters and candidates in a common equilibrium. Turnout depends on both the efficacy of voting, determined by the competitiveness of the election, and the level of campaign outreach. These forces interact: candidates allocate resources to battleground states in an-

ticipation that turnout will respond to strategic effort, while voters adjust their participation based on both intrinsic competitiveness and campaign intensity.

By modeling competitiveness and effort as joint outcomes of strategic behavior, the framework bridges two core literatures that have traditionally been studied in isolation. To my knowledge, it is the first structural model to endogenize both turnout and effort within a single electoral equilibrium.

The model is estimated using county-level data from the 2008 to 2020 presidential elections. Although campaign effort is not directly observed, the model infers it as a latent variable consistent with the candidates' best responses. These inferred allocations match observed patterns of television and digital advertising across battleground states remarkably well. On the voter side, the model's predictions align with reduced-form evidence from a border discontinuity design. This quasi-experimental approach shows that raising intrinsic competitiveness increases turnout by 1.67 percentage points, closely matching the model's implied effect of 1.74 points. This consistency supports the credibility of the estimated efficacy function and the structural separation of competitiveness from mobilization.

The central finding is that voter turnout is not simply motivated by electoral closeness or independently mobilized by campaigns; rather, it is an equilibrium outcome of the interaction between these two forces. Roughly two-thirds of the turnout gap between battleground and non-battleground states is attributable to strategic campaign effort, with the remaining one-third arising from the direct effect of competitiveness on individuals' perceived voting efficacy. While campaign mobilization is shown to be the dominant driver of this gap, the efficacy channel remains a significant and non-negligible factor. Together, these forces account for the entire 6.1 percentage point turnout gap between battleground and non-battleground states.

Beyond accounting for observed turnout patterns, the model also clarifies the nature of campaign mobilization. It reconciles the small and costly marginal effects documented in field experiments with the large aggregate influence of campaign activity in presidential elections. In equilibrium, campaigns invest heavily until the marginal cost of an additional vote exceeds

$250, even though the average cost of mobilization is closer to $150 per vote. This reflects that early voters are relatively easy to mobilize, whereas in saturated campaign environments new voters are much more expensive to reach. Field experiments, which typically measure the effects of mobilization on these later voters, therefore find small impacts at high costs. Yet the cumulative impact is substantial: on average, campaign outreach mobilizes nearly three million additional voters and raises turnout by 4–5 percentage points in pivotal states.

A final counterfactual analysis of public financing reform highlights a fundamental tradeoff. Spending caps would tighten electoral competition and thereby raise turnout through the efficacy channel. Yet this positive effect is offset by the reduction in campaign mobilization. The model predicts that the net result would be a decline in overall participation, underscoring the unintended consequences of policies that restrict campaign investment in voter outreach.

This analysis necessarily abstracts from several complexities, opening avenues for future work. Campaign effort is modeled as targeting only at the state level and assumed to be zero in safe states. The static framework also omits dynamic features such as donor responses, campaign learning, and primary-season investments. Extending the model along these dimensions would enrich its applicability. Incorporating finer geographic targeting, such as DMA-level strategies, could yield more precise predictions of effort. Embedding a fundraising–mobilization tradeoff would provide a richer account of candidate decision-making. Finally, applying the framework to lower-salience elections, including midterms and gubernatorial contests, would test its external validity and broaden its scope.

More broadly, the framework offers a foundation for analyzing electoral participation in settings where changes to the electoral environment influence both voter behavior and campaign strategy. It provides a tractable tool for evaluating a wide range of institutional reforms, from redistricting, which reassigns voters across district lines and reshapes the electoral map, to voting laws that alter participation costs and shift turnout incentives. In both cases, reforms affect not only the intrinsic competitiveness of elections but also the

strategic effort decisions of candidates. A model that jointly endogenizes turnout and effort within a common equilibrium is well suited to capturing these feedbacks and points to several promising directions for future research. Ultimately, the analysis demonstrates that turnout is neither purely "motivated" nor purely "mobilized," but an equilibrium outcome of both.

A Miscellaneous Tables and Figures

Table 16: Estimated Coefficients: Gender and Age

Parameter	Variable	Coefficient	Std. Error	P-Value
μ	Male (2008)	0.533	0.123	0.000
	Male (2012)	0.690	0.118	0.000
	Male (2016)	0.649	0.118	0.000
	Male (2020)	0.425	0.114	0.000
	Age 18–29 (2008)	-0.418	0.075	0.000
	Age 18–29 (2012)	-0.553	0.078	0.000
	Age 18–29 (2016)	-0.645	0.082	0.000
	Age 18–29 (2020)	-0.697	0.081	0.000
	Age 65+ (2008)	-0.308	0.101	0.002
	Age 65+ (2012)	-0.456	0.097	0.000
	Age 65+ (2016)	-0.395	0.093	0.000
	Age 65+ (2020)	-0.460	0.086	0.000

Notes: Coefficients reflect the estimated contribution of each demographic covariate to baseline county-level partisan alignment. Variables are expressed as fractions of county population unless otherwise noted.

Table 17: Estimated Coefficients: Race and Ethnicity

Parameter	Variable	Coefficient	Std. Error	P-Value
μ	White (2008)	1.446	0.220	0.000
	White (2012)	1.451	0.222	0.000
	White (2016)	1.668	0.221	0.000
	White (2020)	1.473	0.167	0.000
	Black (2008)	0.263	0.223	0.238
	Black (2012)	0.252	0.226	0.265
	Black (2016)	0.165	0.222	0.456
	Black (2020)	-0.125	0.168	0.457
	Native American (2008)	0.659	0.234	0.005
	Native American (2012)	0.701	0.238	0.003
	Native American (2016)	0.698	0.231	0.003
	Native American (2020)	0.325	0.174	0.063
	Asian (2008)	1.144	0.298	0.000
	Asian (2012)	1.238	0.300	0.000
	Asian (2016)	1.233	0.290	0.000
	Asian (2020)	1.160	0.232	0.000
	Hispanic (2008)	0.734	0.223	0.001
	Hispanic (2012)	0.737	0.224	0.001
	Hispanic (2016)	0.687	0.220	0.002
	Hispanic (2020)	0.507	0.165	0.002

Notes: See Table 16.

Table 18: Estimated Coefficients: Education

Parameter	Variable	Coefficient	Std. Error	P-Value
μ	High School Only (2008)	-0.217	0.092	0.018
	High School Only (2012)	-0.448	0.099	0.000
	High School Only (2016)	-0.119	0.102	0.244
	High School Only (2020)	-0.194	0.100	0.052
	Some College (2008)	0.303	0.078	0.000
	Some College (2012)	0.087	0.085	0.307
	Some College (2016)	0.350	0.088	0.000
	Some College (2020)	0.110	0.088	0.209
	College Only (2008)	0.220	0.109	0.043
	College Only (2012)	0.123	0.108	0.255
	College Only (2016)	-0.351	0.110	0.001
	College Only (2020)	-0.679	0.103	0.000
	College+ (2008)	-1.976	0.153	0.000
	College+ (2012)	-2.015	0.151	0.000
	College+ (2016)	-2.253	0.149	0.000
	College+ (2020)	-2.525	0.133	0.000

Notes: See Table 16.

Table 19: Estimated Coefficients: Year Effects

Parameter	Variable	Coefficient	Std. Error	P-Value
μ	Year (2008)	-1.233	0.238	0.000
	Year (2012)	-1.228	0.244	0.000
	Year (2016)	-1.314	0.244	0.000
	Year (2020)	-0.674	0.193	0.000

Notes: See Table 16.

B Model Appendix

B.1 Derivation of the Voting Rules

The decision to vote is framed within a standard instrumental voter model. For a given voter i, let:

- u_{iR} be the utility if the Republican candidate (R) wins
- u_{iD} be the utility if the Democratic candidate (D) wins
- $c_i > 0$ be the private cost of voting (e.g., time, effort).

A voter chooses one of three actions: vote for R, vote for D, or abstain. Without loss of generality, assume the voter prefers R to D, so $u_{iR} > u_{iD}$. The choice is therefore between voting for R and abstaining. The voter participates if the expected utility from voting exceeds that from abstaining.

If the voter abstains, their expected utility is the probability-weighted average of the two possible electoral outcomes:

$$\mathbb{E}[U(\text{abstain})] = \Pr(\text{R wins}|\text{i abstains})u_{iR} + \Pr(\text{D wins}|\text{i abstains})u_{iD}. \tag{19}$$

If the voter pays the cost c_i and votes for R, their expected utility becomes:

$$\mathbb{E}[U(\text{vote } R)] = \Pr(\text{R wins}|\text{i votes R})u_{iR} + \Pr(\text{D wins}|\text{i votes R})u_{iD} - c_i. \tag{20}$$

The individual votes for R if the net benefit is positive:

$$\mathbb{E}[U(\text{vote } R)] - \mathbb{E}[U(\text{abstain})] > 0. \tag{21}$$

Note that by definition, the probabilities of winning must sum to one:

$$\Pr(\text{R wins}|\text{i abstains}) + \Pr(\text{D wins}|\text{i abstains}) = 1$$

$$\Pr(\text{R wins}|\text{i votes R}) + \Pr(\text{D wins}|\text{i votes R}) = 1.$$

Substituting equations (19) and (20) into (21) gives the following condition for participation:

$$\begin{aligned}&[\Pr(\text{R wins}|\text{i votes R})u_{iR} + \Pr(\text{D wins}|\text{i votes R})u_{iD} - c_i] \\ &\quad - [\Pr(\text{R wins}|\text{i abstains})u_{iR} + \Pr(\text{D wins}|\text{i abstains})u_{iD}] > 0\end{aligned}$$

Rearranging terms by candidate utility:

$$\begin{aligned}&[\Pr(\text{R wins}|\text{i votes R}) - \Pr(\text{R wins}|\text{i abstains})]\, u_{iR} \\ &\quad + [\Pr(\text{D wins}|\text{i votes R}) - \Pr(\text{D wins}|\text{i abstains})]\, u_{iD} > c_i\end{aligned}$$

Let $\Delta\pi_R$ be the change in the probability of R winning due to voter i's vote. Similarly,

let $\Delta\pi_D$ be the change for D.

$$\Delta\pi_R = \Pr(\text{R wins}|\text{i votes R}) - \Pr(\text{R wins}|\text{i abstains})$$

$$\Delta\pi_D = \Pr(\text{D wins}|\text{i votes R}) - \Pr(\text{D wins}|\text{i abstains})$$

Since probabilities must sum to one, $\Pr(\text{R wins}) + \Pr(\text{D wins}) = 1$, any increase in one candidate's win probability must be matched by an equal decrease in the other's. Therefore, $\Delta\pi_D = -\Delta\pi_R$. This change, $\Delta\pi_R$, is precisely the probability that voter i's vote is pivotal in favor of candidate R. Let's define $P_{\text{pivotal}} \equiv \Delta\pi_R$.

Substituting back into the inequality, we get the final decision rule:

$$P_{\text{pivotal}} \cdot u_{iR} + (-P_{\text{pivotal}}) \cdot u_{iD} > c_i$$

$$P_{\text{pivotal}} \cdot (u_{iR} - u_{iD}) > c_i$$

This is the canonical result: a voter participates when the probability of being pivotal, multiplied by the utility difference, exceeds the cost.

To make this framework empirically tractable and behaviorally plausible, the model deviates from a literal interpretation of P_{pivotal}. Instead of assuming voters calculate precise pivot probabilities (which are vanishingly small in large electorates), I model their subjective belief in their vote's importance as a smooth, continuous function of electoral closeness, denoted $p(\kappa_s)$. The utility stakes, $u_{iR} - u_{iD}$, are captured by the term $|\Delta u_i|$. This leads directly to the voting rule used in section 2.1.3 of the main text:

$$p(\kappa_s) \cdot |\Delta u_i| > c_i$$

B.2 Deriving the Turnout Function

This section derives the county-level turnout rate equations for the Republican and Democratic candidates (equations (4) and (5)), as presented in the main text.

Consider the post-campaign utility differential for individual i in county j_s:

$$\Delta u_{ij_s} = \begin{cases} \min\left\{-m(e_{sD}, e_{sR}) + \Delta\tilde{u}_{ij_s} + \zeta_{j_s},\ 0\right\}, & \text{if } \Delta\tilde{u}_{ij_s} < 0, \\ \max\left\{m(e_{sR}, e_{sD}) + \Delta\tilde{u}_{ij_s} - \zeta_{j_s},\ 0\right\}, & \text{if } \Delta\tilde{u}_{ij_s} \geq 0, \end{cases}$$

where the baseline utility differential is

$$\Delta\tilde{u}_{ij_s} = \mu_{j_s} - \eta_{j_s} - \delta_s - \epsilon_{ij_s}.$$

Republican Turnout

For an R-leaning voter ($\Delta\tilde{u}_{ij_s} \geq 0$), turnout occurs if

$$p(\kappa_s) \cdot \Delta u_{ij_s} > c_{j_s}.$$

Substituting the definition of Δu_{ij_s}:

$$p(\kappa_s) \cdot \max\left\{m(e_{sR}, e_{sD}) + \Delta\tilde{u}_{ij_s} - \zeta_{j_s},\ 0\right\} > c_{j_s}.$$

Note that $c_{j_s} > 0$ and $p(\kappa_s) > 0$. The condition $\max\{A, 0\} > B$ with $B > 0$ is equivalent to $A > B$. The turnout condition therefore simplifies to

$$m(e_{sR}, e_{sD}) + \Delta\tilde{u}_{ij_s} - \zeta_{j_s} > \frac{c_{j_s}}{p(\kappa_s)}.$$

Substituting $\Delta\tilde{u}_{ij_s}$ and rearranging:

$$\epsilon_{ij_s} < m(e_{sR}, e_{sD}) + \mu_{j_s} - \frac{c_{j_s}}{p(\kappa_s)} - \eta_{j_s} - \delta_s - \zeta_{j_s}.$$

Therefore an individual votes for the Republican candidate if this condition holds, and $\epsilon_{ij_s} < \mu_{j_s} - \eta_{j_s} - \delta_s$ (i.e., the individual is R-leaning). This gives us the joint probability of voting for R:

$$\Pr(\text{Vote } R) = \Pr\left(\epsilon_{ij_s} < m(e_{sR}, e_{sD}) + \mu_{j_s} - \frac{c_{j_s}}{p(\kappa_s)} - \eta_{j_s} - \delta_s - \zeta_{j_s},\ \epsilon_{ij_s} < \mu_{j_s} - \eta_{j_s} - \delta_s\right)$$

Assume that not all R-leaning voters in a given county turn out. With thousands of voters per county, this is a mild restriction, and it is verified post-estimation. Under this assumption, the first condition is the more restrictive one, implying:

$$m(e_{sR}, e_{sD}) - \tfrac{c_{j_s}}{p(\kappa_s)} - \zeta_{j_s} < 0$$

Therefore, using the fact that $\epsilon_{ij_s} \sim N(0,1)$ with CDF $H(\cdot)$, the Republican turnout rate in county j_s is

$$\sigma_{j_s R} = H\left(m(e_{sR}, e_{sD}) + \mu_{j_s} - \frac{c_{j_s}}{p(\kappa_s)} - \eta_{j_s} - \delta_s - \zeta_{j_s}\right).$$

Democratic Turnout

For a D–leaning voter ($\Delta\tilde{u}_{ij_s} < 0$), turnout occurs if

$$p(\kappa_s)\,|\Delta u_{ij_s}| > c_{j_s}.$$

Since

$$\Delta u_{ij_s} = \min\{-m(e_{sD}, e_{sR}) + \Delta\tilde{u}_{ij_s} + \zeta_{j_s},\, 0\} \leq 0,$$

we have

$$|\Delta u_{ij_s}| = \max\{m(e_{sD}, e_{sR}) - \Delta\tilde{u}_{ij_s} - \zeta_{j_s},\, 0\}\,.$$

Because $c_{j_s}/p(\kappa_s) > 0$, the turnout condition is equivalent to

$$m(e_{sD}, e_{sR}) - \Delta\tilde{u}_{ij_s} - \zeta_{j_s} > \frac{c_{j_s}}{p(\kappa_s)}.$$

Substituting $\Delta\tilde{u}_{ij_s} = \mu_{j_s} - \eta_{j_s} - \delta_s - \epsilon_{ij_s}$:

$$p(\kappa_s) \cdot [m(e_{sD}, e_{sR}) - \mu_{j_s} + \eta_{j_s} + \delta_s + \epsilon_{ij_s} - \zeta_{j_s}] > c_{j_s}.$$

Dividing through by $p(\kappa_s)$ and rearranging:

$$\epsilon_{ij_s} > \frac{c_{j_s}}{p(\kappa_s)} - m(e_{sD}, e_{sR}) + \mu_{j_s} - \eta_{j_s} - \delta_s + \zeta_{j_s}.$$

Again, assuming that not all D-leaning voters turn out, we can ignore the redundant condition $\epsilon_{ij_s} > \mu_{j_s} - \eta_{j_s} - \delta_s$. Because $\epsilon_{ij_s} \sim N(0,1)$ with CDF $H(\cdot)$, the probability that this condition holds is

$$1 - H\left(\frac{c_{j_s}}{p(\kappa_s)} - m(e_{sD}, e_{sR}) + \mu_{j_s} - \eta_{j_s} - \delta_s + \zeta_{j_s}\right).$$

Using $1 - H(x) = H(-x)$, the Democratic turnout rate in county j_s is:

$$\sigma_{j_sD} = H\left(m(e_{sD}, e_{sR}) - \mu_{j_s} - \frac{c_{j_s}}{p(\kappa_s)} + \eta_{j_s} + \delta_s - \zeta_{j_s}\right).$$

B.3 Central Limit Theorem Validation

The candidate's objective function relies on a normal approximation for the distribution of total electoral votes. To validate this key simplifying assumption, I conduct a Monte Carlo simulation for each election cycle. Using the estimated state-win probabilities ($\tilde{\pi}_s$) from the

structural model, I simulate the election outcome 1,000,000 times, generating an empirical distribution of electoral votes for the Democratic candidate. Figure 11 compares the histograms of these simulated outcomes against the theoretical normal distribution predicted by the model. The close alignment for each year confirms that the normal approximation is accurate and robust, providing a solid foundation for the analysis of the candidates' equilibrium strategies.

B.4 Central Limit Theorem Validation

The candidate's objective function (Equation 6) relies on a normal approximation for the distribution of total electoral votes. To validate this key simplifying assumption, I conduct a Monte Carlo simulation for each election cycle. Using the estimated state-win probabilities ($\tilde{\pi}_s$) from the structural model, I simulate the election outcome 100,000 times, generating an empirical distribution of electoral votes for the Democratic candidate. Figure 11 compares the histograms of these simulated outcomes against the theoretical normal distribution predicted by the model. The close alignment for each year confirms that the normal approximation is accurate and robust, providing a solid foundation for the analysis of the candidates' equilibrium strategies.

B.5 Accuracy of the County-Shock Approximation

This section quantifies the approximation error from omitting county-level shocks, η_{j_s} and ζ_{j_s}, when computing the probability that the Democrat wins state s. I conduct $N_{\text{sim}} = 250$ Monte Carlo simulations. In each simulation, I draw $N_{\text{trial}} = 10{,}000$ independent realizations of the shocks to estimate both the full and approximate win probabilities.

Each simulation randomly draws a number of precincts for state s, campaign effort levels for both parties, and the structural model parameters. Then, for each of the N_{trial} draws, I sample:

$$\eta_{j_s} \sim \mathcal{N}(0, \sigma_j^2), \qquad \zeta_{j_s} \sim \mathcal{N}(0, \sigma_j^2), \qquad \delta_s \sim \mathcal{N}(0, \sigma_s^2),$$

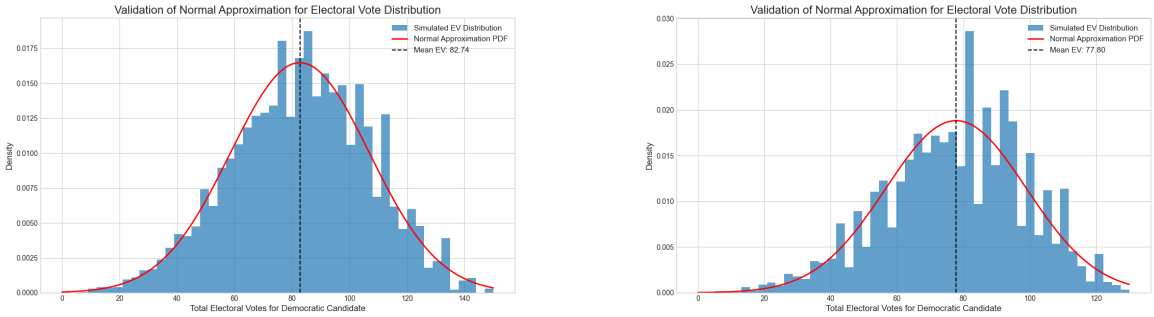

(a) 2008 Election

(b) 2012 Election

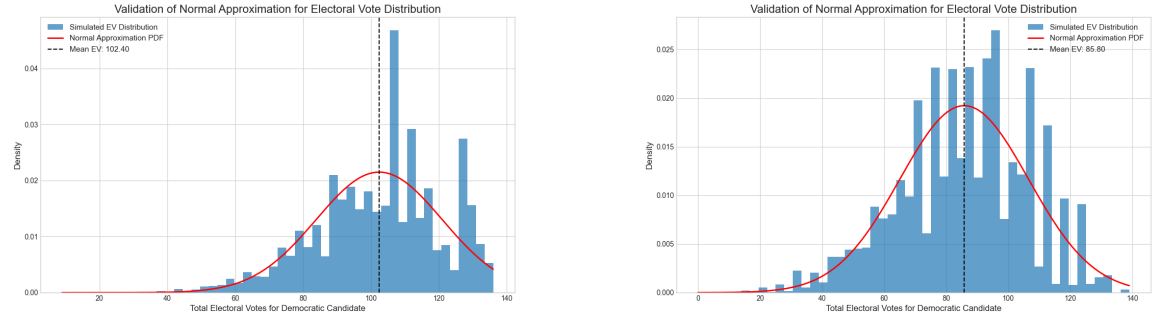

(c) 2016 Election

(d) 2020 Election

Figure 11: Validation of the Normal Approximation for the Electoral Vote Distribution. Each panel shows the histogram of total electoral votes for the Democratic candidate from 1,000,000 simulations, using the estimated model parameters for that election year. The solid red line is the corresponding Normal PDF, and the dashed black line indicates the mean of the simulated distribution.

and compute two probabilities:

$$\pi_D^{\text{full}} = \Pr\big[\sigma_{sD}(\eta, \zeta, \delta) > \sigma_{sR}(\eta, \zeta, \delta)\big],$$

$$\pi_D^{\text{approx}} = \Pr\big[\sigma_{sD}(0, 0, \delta) > \sigma_{sR}(0, 0, \delta)\big],$$

where $\sigma_{sD}(\cdot)$ and $\sigma_{sR}(\cdot)$ denote the state-level turnout functions for Democrats and Republicans, respectively, defined in Equations (4) and (5). Each probability is estimated as the fraction of draws in which the Democrat's simulated vote share exceeds the Republican's.

Let $\Delta_i = |\pi_D^{\text{full}} - \pi_D^{\text{approx}}|$ denote the absolute error in simulation i. I report the mean and median absolute error across simulations, along with the standard deviation, minimum,

maximum, and percentiles of the absolute error distribution:

$$\begin{aligned} \text{Mean}(\Delta) &= \frac{1}{N_{\text{sim}}} \sum_{i=1}^{N_{\text{sim}}} \Delta_i, \\ \text{Median}(\Delta) &= \text{median}(\Delta_1, \ldots, \Delta_{N_{\text{sim}}}). \end{aligned}$$

The average absolute error is 0.0099, and the median is 0.0062. Thus, the full and approximate probabilities differ by less than one percentage point in expectation, and by just over a half of a percentage point in the median case.

B.6 Numerical Verification of the Uniqueness of $\hat{\delta}_s$

This section verifies that the threshold value $\hat{\delta}_s$ defined in 9 is unique for a broad range of parameter values and campaign effort profiles.[12]

Each simulation randomly draws a number of precincts for state s, campaign effort levels for both parties, and the structural model parameters. For each of the $N_{\text{sim}} = 1{,}000$ Monte Carlo simulations, I attempt to solve for $(\hat{\delta}_s, \sigma_{s,D}, \sigma_{s,R})$ from 9, using $N_{\text{trial}} = 100$ *independent* random initial guesses. The root-finding routine is a damped Newton method implemented in `JAX`, with the Jacobian computed via automatic differentiation.

Let $\boldsymbol{\delta}^{(r,t)} = (\hat{\delta}_s, \sigma_{s,D}, \sigma_{s,R})^{(r,t)}$ denote the solution in simulation r (out of N_{sim}) and trial t (out of N_{trial}). Trials that fail to converge or that return the degenerate boundary solution $(x, 0, 0)$ are discarded.

To assess dispersion in solutions across initializations, I compute the coordinate-wise variance:

$$\text{Var}_t(\boldsymbol{\delta}^{(r,t)}) = \big(\text{Var}_t(\hat{\delta}_s^{(r,t)}),\ \text{Var}_t(\sigma_{s,D}^{(r,t)}),\ \text{Var}_t(\sigma_{s,R}^{(r,t)})\big),$$

[12] In principle, two distinct fixed points may exist: an interior solution and a degenerate boundary solution of the form $(x, 0, 0)$, where x satisfies 9 under $\sigma_{s,D} = \sigma_{s,R} = 0$. I exclude such degenerate solutions from the analysis, as they do not correspond to meaningful equilibria.

where $\mathrm{Var}_t(\cdot)$ denotes variance across trials t within a given simulation r. If multiple fixed points existed, at least one coordinate of this vector would be strictly positive.

Across the 1,000 simulations, the largest coordinate-wise variance is given in Table 20.

Table 20: Dispersion of $(\hat{\delta}_s, \sigma_{s,D}, \sigma_{s,R})$ across Random Initializations

	$\hat{\delta}_s$	$\sigma_{s,D}$	$\sigma_{s,R}$
$\max_r \mathrm{Var}_t(\cdot)$	1.37×10^{-23}	5.05×10^{-22}	5.28×10^{-22}

Notes: Values report the largest coordinate-wise variance observed over $N_{\text{sim}} = 1{,}000$ simulations and $N_{\text{trial}} = 100$ random starting points per simulation. The vanishing dispersion indicates convergence to a *single* fixed point in every replication.

In every simulation, one of three outcomes occurred: (i) convergence to a unique interior fixed point, (ii) convergence to the boundary solution $(x, 0, 0)$, or (iii) failure to converge. Since a valid interior solution was recovered in every parameter draw, and coordinate-wise dispersion is vanishingly small, I conclude that the solution to 9 is globally unique whenever an interior solution exists. Hence, the smoothed win probability, $\tilde{\pi}_s(e_{s,D}, e_{s,R}) = 1 - F(\hat{\delta}_s)$, is well defined.

C Data Appendix

C.1 Summary statistics

Table 21: Summary Statistics for Covariates Used in the Partisan-Bias Parameter (Battleground States)

Variable	Count	Mean	Std. Dev.	Min	25%	50%	75%	Max
Male	3,174	0.501	0.026	0.374	0.488	0.496	0.505	0.790
Age 18–29	3,174	0.147	0.047	0.034	0.123	0.137	0.156	0.591
Age 65+	3,174	0.162	0.044	0.028	0.134	0.159	0.185	0.514
White	3,174	0.789	0.176	0.086	0.680	0.852	0.932	0.998
Black	3,174	0.103	0.143	0.000	0.008	0.033	0.145	0.791
Native American	3,174	0.008	0.037	0.000	0.001	0.002	0.004	0.796
Asian	3,174	0.013	0.018	0.000	0.003	0.006	0.014	0.189
Hispanic	3,174	0.066	0.082	0.000	0.019	0.039	0.078	0.831
High School Only	3,174	0.352	0.077	0.055	0.304	0.354	0.405	0.556
Some College	3,174	0.297	0.048	0.114	0.265	0.300	0.330	0.455
College Only	3,174	0.142	0.059	0.030	0.098	0.130	0.170	0.480
College+	3,174	0.077	0.046	0.007	0.047	0.064	0.094	0.437

Notes: All variables are expressed as population shares unless otherwise noted. Statistics are based on county-level data from battleground states where effort is endogenously allocated in equilibrium.

Table 22: Summary Statistics for Covariates Used in the Partisan-Bias Parameter (Non-Battleground States)

Variable	Count	Mean	Std. Dev.	Min	25%	50%	75%	Max
Male	9,322	0.501	0.023	0.405	0.489	0.497	0.507	0.764
Age 18–29	9,322	0.147	0.041	0.031	0.125	0.141	0.159	0.554
Age 65+	9,322	0.157	0.041	0.031	0.129	0.154	0.181	0.402
White	9,322	0.763	0.207	0.007	0.647	0.839	0.927	1.000
Black	9,322	0.084	0.144	0.000	0.005	0.018	0.085	0.874
Native American	9,322	0.021	0.081	0.000	0.001	0.003	0.007	0.910
Asian	9,322	0.013	0.030	0.000	0.002	0.005	0.011	0.522
Hispanic	9,322	0.097	0.148	0.000	0.020	0.038	0.098	0.991
High School Only	9,322	0.345	0.070	0.071	0.300	0.349	0.396	0.557
Some College	9,322	0.301	0.055	0.111	0.265	0.301	0.338	0.506
College Only	9,322	0.136	0.055	0.019	0.095	0.126	0.166	0.457
College+	9,322	0.072	0.042	0.000	0.045	0.060	0.086	0.483

Notes: All variables are expressed as fractions of county population unless otherwise noted. Observations cover counties in non-battleground states where campaign effort is set to zero in equilibrium.

Table 23: Summary Statistics for Cost of Voting Index by Urban/Rural Status and State Type

	Urban Counties		Rural Counties	
Statistic	BG	Non-BG	BG	Non-BG
Count	231	375	2880	8324
Mean	0.584	0.593	0.540	0.521
Std. Dev.	0.042	0.050	0.054	0.064
Min	0.471	0.491	0.334	0.000
25th Pctl	0.559	0.563	0.505	0.482
Median	0.590	0.587	0.536	0.519
75th Pctl	0.609	0.613	0.570	0.558
Max	0.760	0.976	0.887	1.000

Notes: The cost of voting index reflects percentile-transformed congestion measures interacted with urban status. Urban counties are defined as those with more than 350 residents per square kilometer. BG = Battleground states (receive campaign effort); Non-BG = Non-battleground states (no campaign effort). All statistics are computed at the county-election level. Values are at the county-election level.

Table 24: Summary Statistics for Election-Law Indices by State Type

	Voter ID Index		Poll Hours Index		Reg. Deadline Index	
Statistic	BG	Non-BG	BG	Non-BG	BG	Non-BG
Count	40	160	40	160	40	160
Mean	0.300	0.273	0.377	0.380	0.519	0.563
Std. Dev.	0.345	0.317	0.142	0.207	0.463	0.406
Min	0.000	0.000	0.000	0.000	0.000	0.000
25th Pctl	0.000	0.000	0.333	0.333	0.000	0.000
Median	0.250	0.250	0.333	0.417	0.717	0.700
75th Pctl	0.500	0.500	0.500	0.500	0.967	0.967
Max	1.000	1.000	0.667	1.000	1.000	1.000

Notes: All indices are scaled to the unit interval. Higher values reflect more restrictive voting policies. BG = Battleground states (receive campaign effort); Non-BG = Non-battleground states (no campaign effort). Values is at the state-year level.

Table 25: Summary Statistics for Cost of Voting Covariates by State Type

	Fraction Employed	
Statistic	BG	Non-BG
Count	3,174	9,059
Mean	0.598	0.589
Std. Dev.	0.080	0.079
Min	0.136	0.000
25th Pctl	0.551	0.541
Median	0.606	0.597
75th Pctl	0.654	0.645
Max	0.855	0.831

Notes: Variable normalized to the unit interval using empirical percentiles. BG = Battleground states (receive campaign effort); Non-BG = Non-battleground states (no campaign effort). Statistics are based on county-election level observations.

C.2 Polling-place congestion

I proxy queues at the polls with a crowding index that scales the voting-age population by the number of in-person polling locations on Election Day. Let PP_{jt} denote the number of polling places and VAP_{jt} the voting-age population in county j during election year t, defined as the number of residents aged 18 and older. The congestion index is defined as

$$Congest_{jt} = \log\left(\frac{VAP_{jt}}{\widehat{PP}_{jt}}\right),$$

where $\widehat{PP}_{jt}$ is the observed number of polling sites, or a predicted value when the data is missing.

Approximately 19% of county-election observations are missing polling place data. I im-

pute these missing values using a Gradient Boosting model trained on county-year covariates, including a linear time trend, log voting-age population, demographic shares (age, gender, race, education, employment), and state fixed effects. I first log-transform the number of polling places to reduce skewness and ensure positive predictions. Model hyperparameters are selected via five-fold cross-validated grid search over tree depth, regularization, learning rate, and number of iterations. The model is trained on observed data from the 2008-2020 cycles. The final model achieves an out-of-sample $R^2 = 0.838$.

C.3 Battleground state classification

To illustrate the sharp concentration of campaign resources, Figure 12 plots the marginal increase in total advertising expenditures from adding each successive state to the battleground set, ranked in descending order of combined Democratic and Republican spending between August 1 and Election Day. Across all years, the first few states produce large jumps in total spending, but the marginal gain falls rapidly. By the tenth state, additional states contribute negligibly to overall expenditures, confirming that campaign activity is overwhelmingly focused on a small set of states.

Figure 12: Percent Increase from Adding States to the Battleground Classification

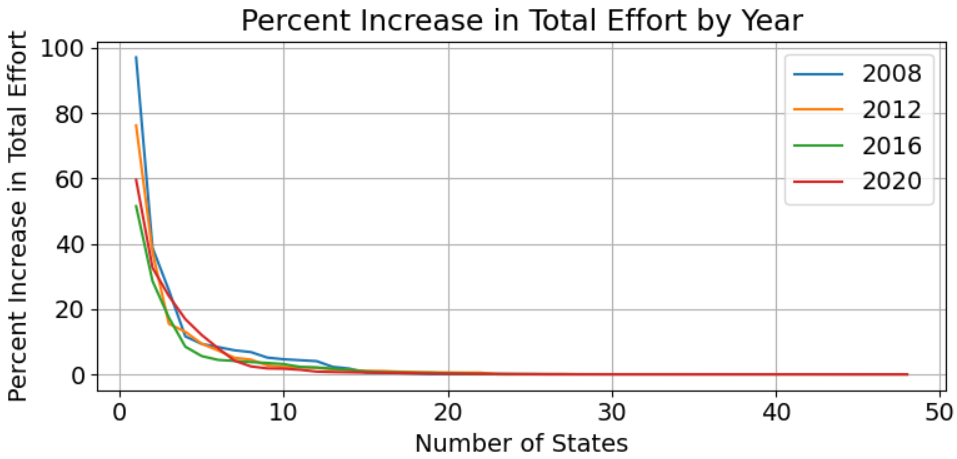

Notes: Each curve shows the percent increase in total advertising expenditures from adding the next state to the battleground set, with states ranked in descending order of combined Democratic and Republican television spending between August 1 and Election Day. Data are from the Wesleyan Media Project for the 2008, 2012, 2016, and 2020 presidential elections.

Tables 26 and 27 summarize the distribution of observed television advertising expenditures across states. In each election year, the top ten states account for more than 86% of total spending, with this share rising to 92% in 2020. This sharp concentration motivates the definition of battleground states used in the model.

Among the remaining 40 states, most receive negligible effort: the median share is effectively zero in every year, and even the 75th percentile remains well below 1% in all cycles. These patterns support the assumption that campaign effort is zero in non-battleground states. While this imposes a discrete cutoff, it closely mirrors the observed data and substantially simplifies the model's strategic problem without distorting the distribution of effort.

Table 26: Summary Statistics for Share of Total Effort in Non-Battleground States

Statistic	2008	2012	2016	2020
Mean	0.0045	0.0036	0.0029	0.0022
Std. Dev.	0.0110	0.0060	0.0053	0.0041
Min	0.0000	0.0000	0.0000	0.0000
25th pct.	0.0000	0.0000	0.0001	0.0000
Median	0.0000	0.0000	0.0002	0.0000
75th pct.	0.0010	0.0071	0.0030	0.0022
Max	0.0448	0.0196	0.0221	0.0163

Table 27: Share of Total Campaign Effort in Top 10 Battleground States

Year	2008	2012	2016	2020
Share	0.830	0.862	0.862	0.924

D Constructing Campaign Budget Shares and Total Effort

D.1 Data Sources

Campaign–finance information comes from two datasets:

- *Television advertising.* Gross state-level outlays on presidential television ads are provided by the Wesleyan Media Project. These figures form the variable $\text{TV}_{s,p}$ discussed in Section 4.4.
- *Operating expenditures.* Itemized operating-expenditure files released by the Federal Election Commission (FEC) record every payment made by candidate committees, including transaction date, amount, and free-text purpose description.

D.2 Filtering Operating Expenditures

The raw FEC files contain many transactions unrelated to voter mobilization. The following rules are applied to retain only plausible mobilization outlays:

1. **General-election focus:** Keep entries tagged as general or general–primary spending.

2. **Candidate committees:** Restrict to disbursements by the principal presidential committees of each major party.

3. **Purpose description cleaning:** Convert purpose strings to lower-case and harmonize common variants (e.g. "on-line" $\rightarrow$ "online").

4. **Positive keyword match:** Retain only transactions whose purpose description matches one of the five predefined mobilization categories (media, online, print, telemarketing, travel) based on regular expressions.

D.3 Classifying Mobilization Channels

Every retained transaction is assigned to one of five mutually exclusive mobilization categories using keyword patterns:

Category	Matched keywords in purpose description
`media`	media, tv, broadcast
`online`	online, digital, facebook, google, youtube, twitter, instagram, snapchat, web, internet
`print`	print, post, mail, leaflet
`telemarketing`	telemarketing, phone, text, sms
`travel`	travel, event, rally, airfare, hotel

Ambiguous strings are resolved by a priority order media $\succ$ online $\succ$ print $\succ$ telemarketing $\succ$ travel, ensuring each transaction appears exactly once.

E Identification Appendix

E.1 Identification via Monte Carlo Simulation

For the first validation exercise, I use the observed county- and state-level covariates from the empirical application, $\left(X_{j_s}^{\mu}, X_{j_s}^{c}\right)$, to generate $R = 100$ synthetic elections. In each replication, I draw a fresh vector of coefficients

$$(\beta_{\mu}, \beta_{c}, \beta_{\alpha_1}, \beta_{\alpha_2}, \beta_{\theta}, \beta_{\eta}, \beta_{\delta}),$$

compute equilibrium campaign efforts $\left(e_{s,D}^{(r)}, e_{s,R}^{(r)}\right)$, draw county- and state-level shocks, and simulate turnout. I then re-estimate the model on each simulated dataset using the same likelihood function and optimization routine as in the baseline estimation. To economize on computation time, I use a reduced covariate set and fix $\gamma = 2$ and $\psi = 0$. Replications that fail to converge or do not yield an equilibrium profile are excluded, as this issue does not arise in the empirical application. Table 28 reports the results. Mean estimation errors are centered at zero, indicating unbiasedness, and mean squared errors are small, indicating high precision. This confirms that the structural parameters can be reliably recovered.

Table 28: Parameter-recovery diagnostics across Monte-Carlo replications

Parameter	MSE	RMSE	25th pctl	75th pctl
Male (18–29)	1.89×10^{-4}	0.0138	−0.0078	0.0077
Female (65–79)	5.11×10^{-4}	0.0226	−0.0127	0.0147
White	2.12×10^{-5}	0.0046	−0.0029	0.0021
Black	2.92×10^{-5}	0.0054	−0.0041	0.0028
Hispanic	2.37×10^{-5}	0.0049	−0.0026	0.0028
High school only	1.04×10^{-4}	0.0102	−0.0052	0.0063
Some college	5.19×10^{-5}	0.0072	−0.0022	0.0060
College only	1.54×10^{-4}	0.0124	−0.0063	0.0078
Employed	8.91×10^{-6}	0.0030	−0.0013	0.0019
Voter ID Index	8.02×10^{-7}	0.00090	−0.00041	0.00049
Cost constant	3.54×10^{-6}	0.00188	−0.00103	0.00085
α_1	2.59×10^{-2}	0.1609	−0.0265	0.0204
α_2	2.46×10^{-1}	0.4956	−0.0581	0.0863
θ	5.06×10^{-5}	0.0071	−0.0040	0.0037
σ_c	5.92×10^{-5}	0.0077	0.00033	0.0080
σ_s	6.96×10^{-3}	0.0834	−0.0259	0.0623

Notes: MSE is the mean squared estimation error across Monte-Carlo replications, and RMSE is its square root. The final two columns report the 25th and 75th percentiles of the estimation error distribution for each parameter.

E.2 Identification on estimated coefficients

To assess local identification and numerical stability, I conduct a likelihood sensitivity analysis around the estimated parameter vector $\hat{\beta}$. For each element β_k, I generate a grid of values in a neighborhood around $\hat{\beta}_k$, holding all other elements fixed. At each grid point, I re-evaluate the full-sample log-likelihood function and its gradient.

Formally, let $\hat{\beta} \in R^K$ denote the estimated parameter vector, and fix a grid of shocks $\{\delta_m\}_{m=1}^M$. For each index $k = 1, \ldots, K$, I define a perturbed parameter vector $\beta^{(k,m)}$ such that

$$\beta_j^{(k,m)} = \begin{cases} \hat{\beta}_j + \delta_m, & \text{if } j = k, \\ \hat{\beta}_j, & \text{otherwise.} \end{cases}$$

For each perturbed vector, I compute the log-likelihood $\ell(\beta^{(k,m)})$ and record the results. This procedure yields a series of univariate likelihood profiles centered at $\hat{\beta}_k$, which allow visual inspection of local curvature and potential flat regions in the likelihood surface. In practice, I fix δ_m to be between -1 and 1, with a total of 20 grid points for each β_k.

Figures 13 - 16 plot these likelihood profiles for each β_k. The results show that the likelihood is locally well-behaved and concave in the neighborhood of each coefficient. No flat regions or multimodalities are detected, providing reassurance that the likelihood-based estimator is locally identified and numerically stable. The only exception are the α_1 and the α_2 parameters, which govern the perceived efficacy function. These parameters exhibit a flatter likelihood profile. However, the likelihood still exhibits a midpoint around the true value, suggesting that the model is still locally identified, albeit with less precision for these parameters.

F Results Appendix

F.1 Calculating Elasticity of Turnout with Respect to Campaign Effort

To quantify marginal responsiveness, I exploit the equilibrium conditions for county-level turnout, defined in equations (4) and (5). In equilibrium, the following system must be satisfied for each county j_s in state s and each party $p \in \{D, R\}$:

Figure 13: Log-likelihood profiles around each β_k (Graph 1)

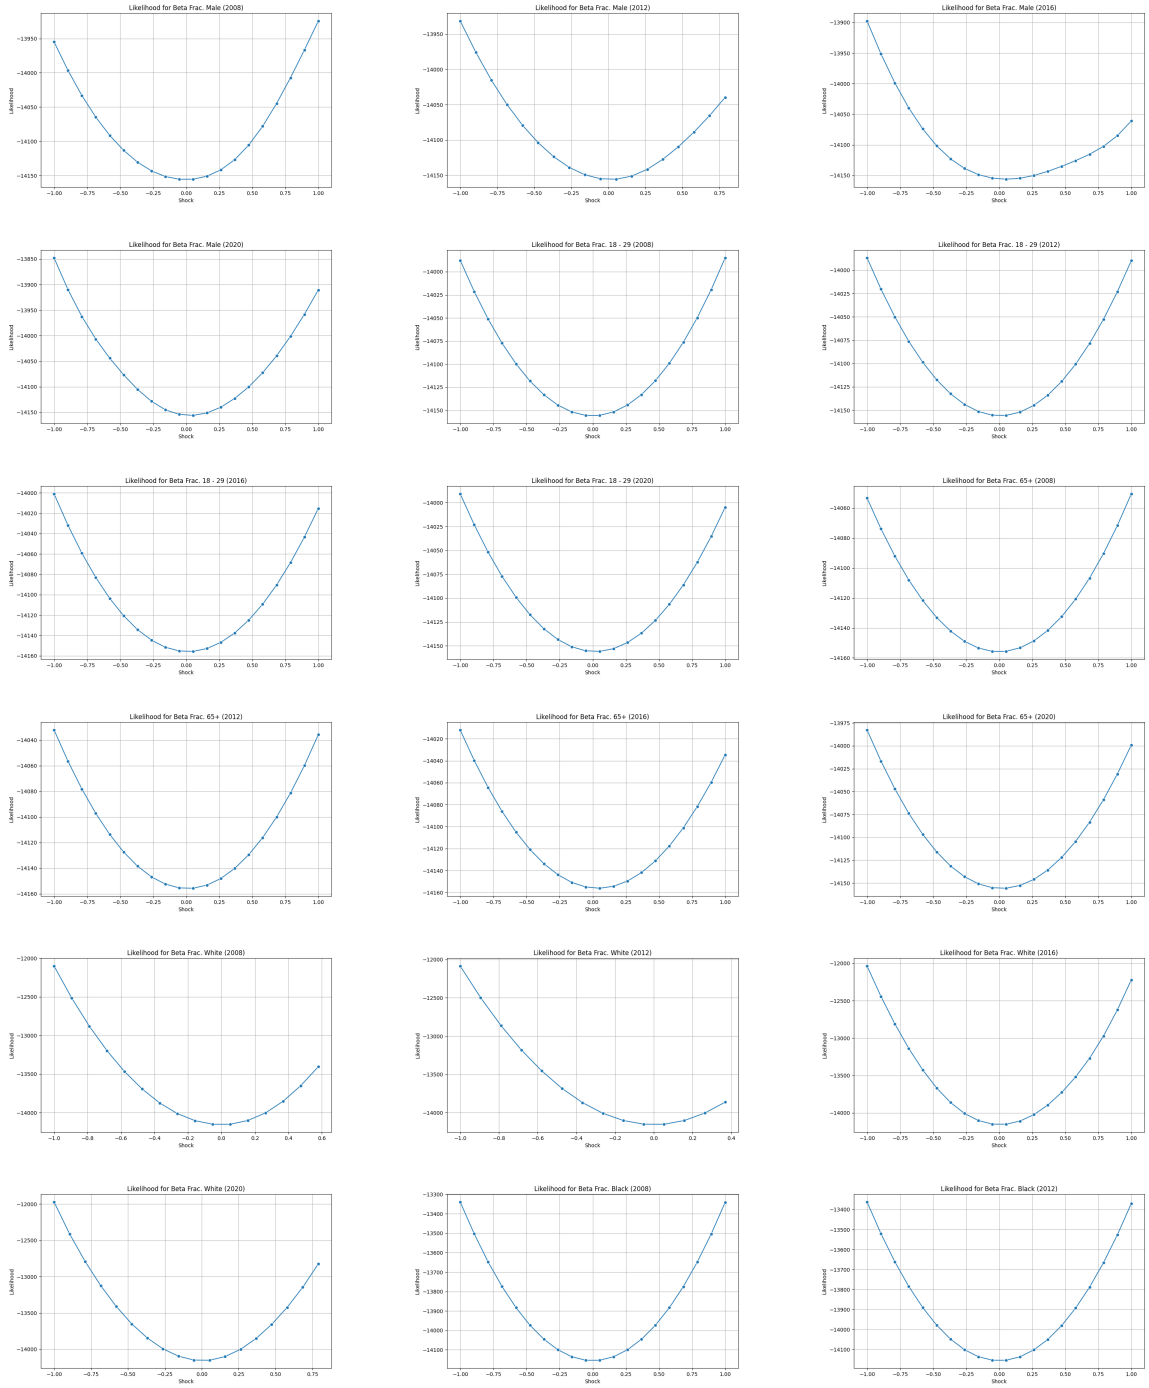

Log-likelihood evaluated around each β_k by applying additive shocks. Each panel holds all other coefficients fixed and perturbs only the k-th element.

Figure 14: Log-likelihood profiles around each β_k (Graph 2)

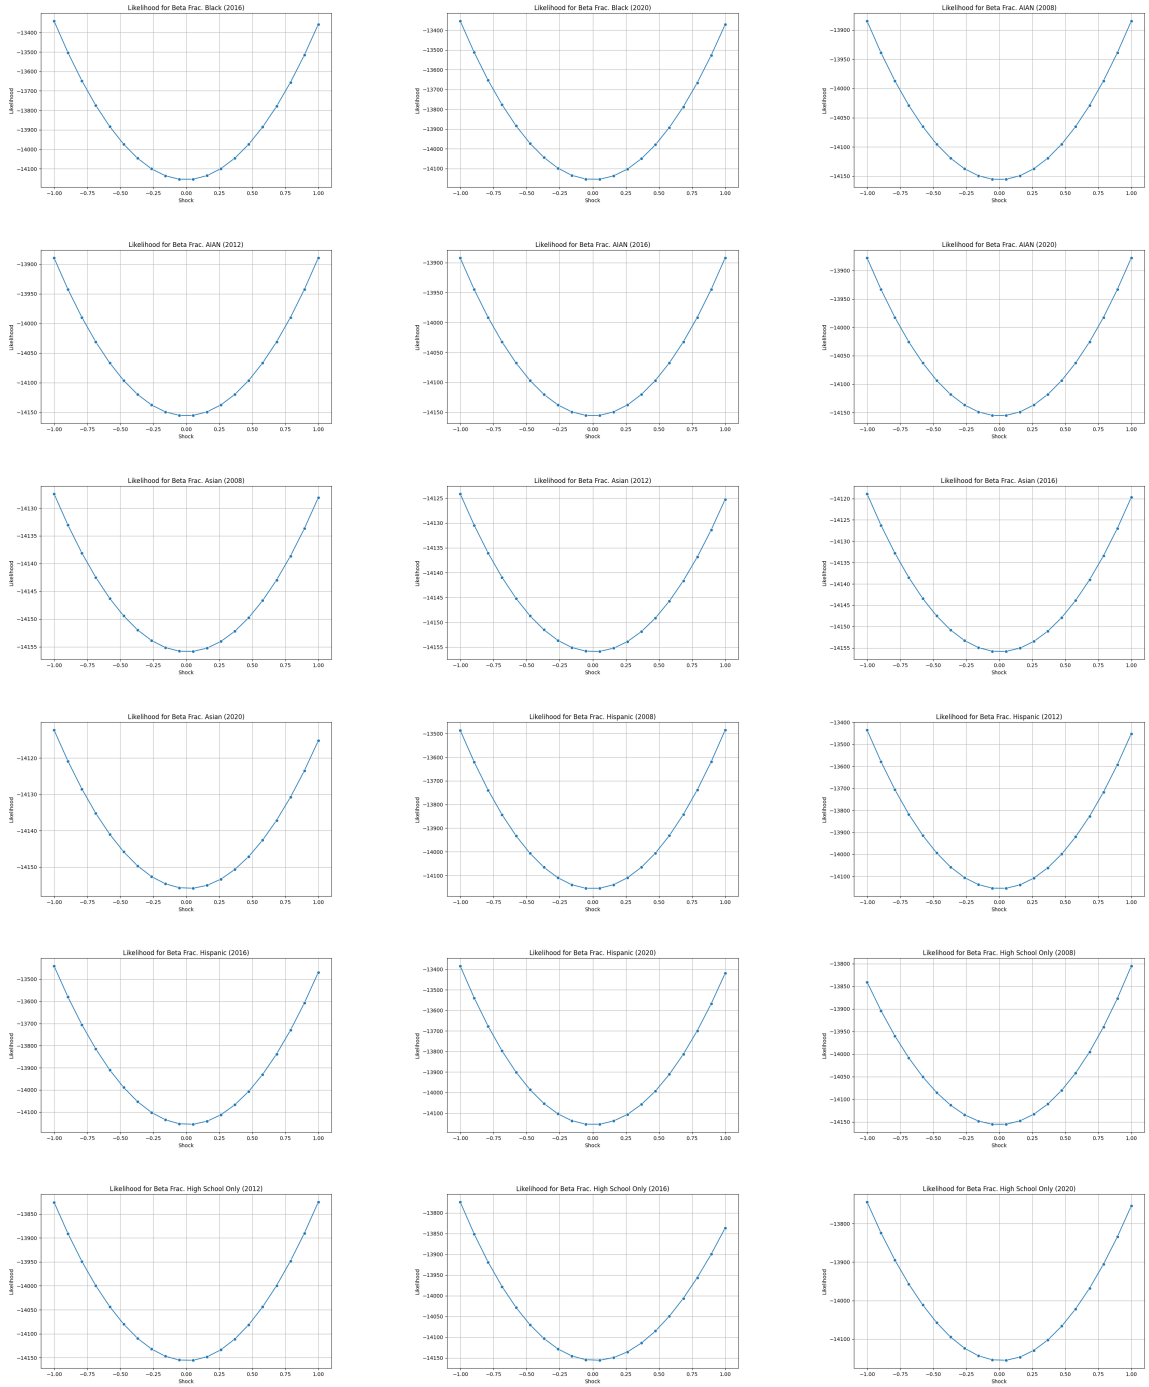

Log-likelihood evaluated around each β_k by applying additive shocks. Each panel holds all other coefficients fixed and perturbs only the k-th element.

Figure 15: Log-likelihood profiles around each β_k (Graph 3)

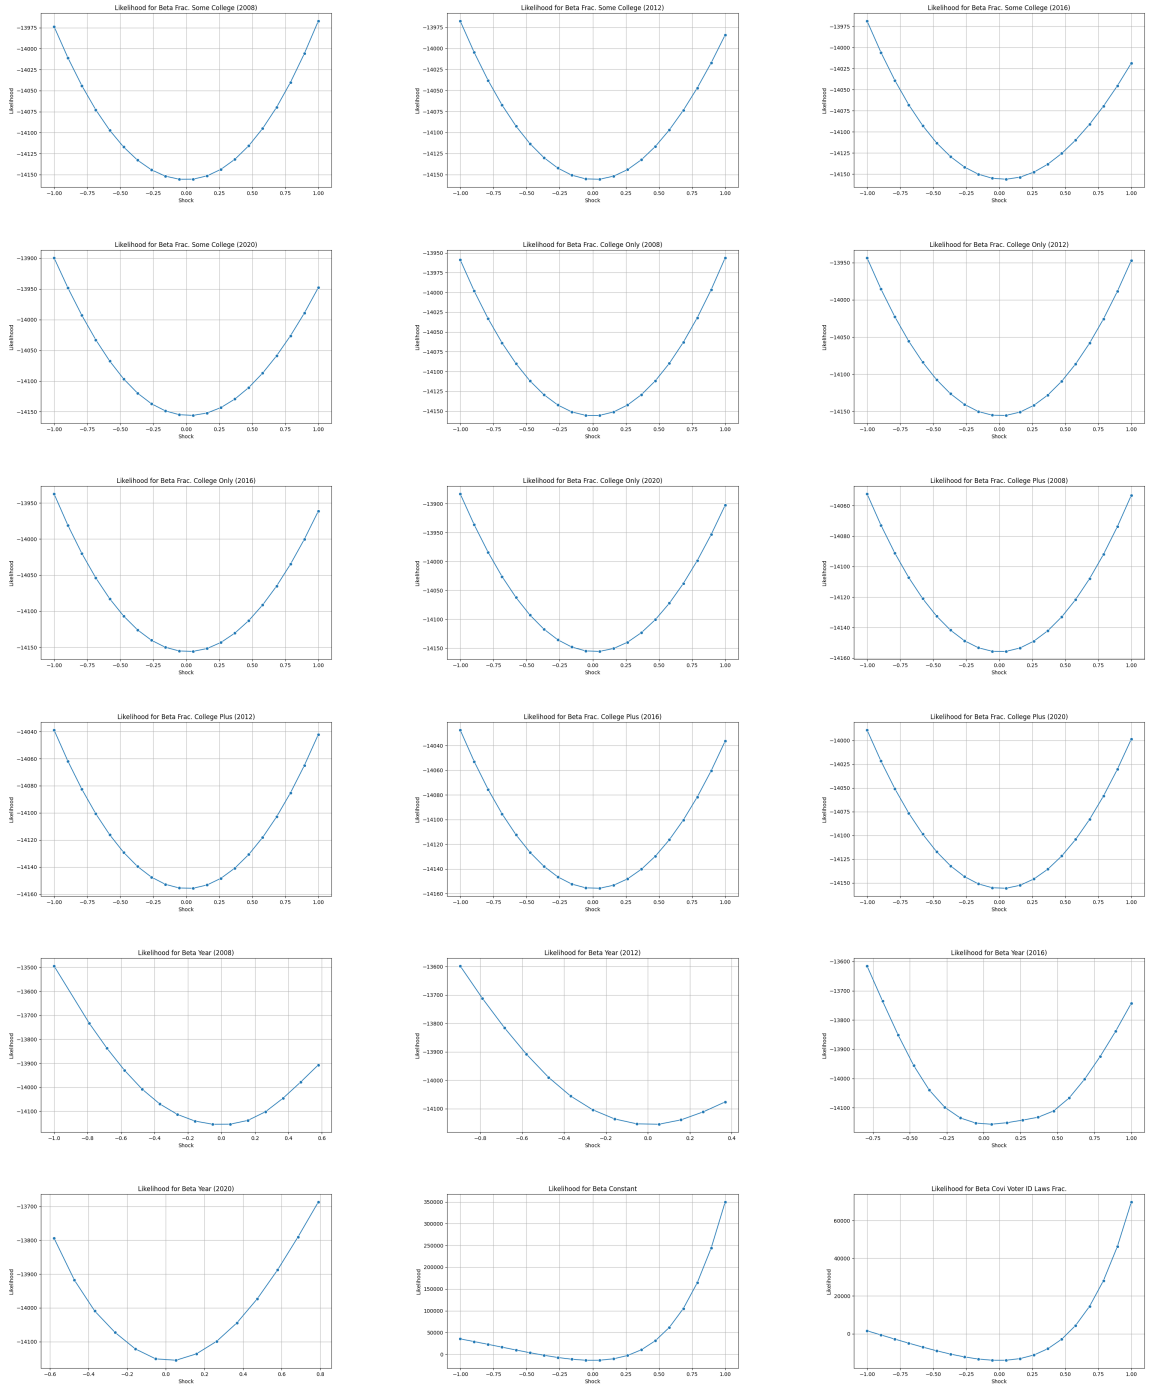

Log-likelihood evaluated around each β_k by applying additive shocks. Each panel holds all other coefficients fixed and perturbs only the k-th element.

Figure 16: Log-likelihood profiles around each β_k (Graph 4)

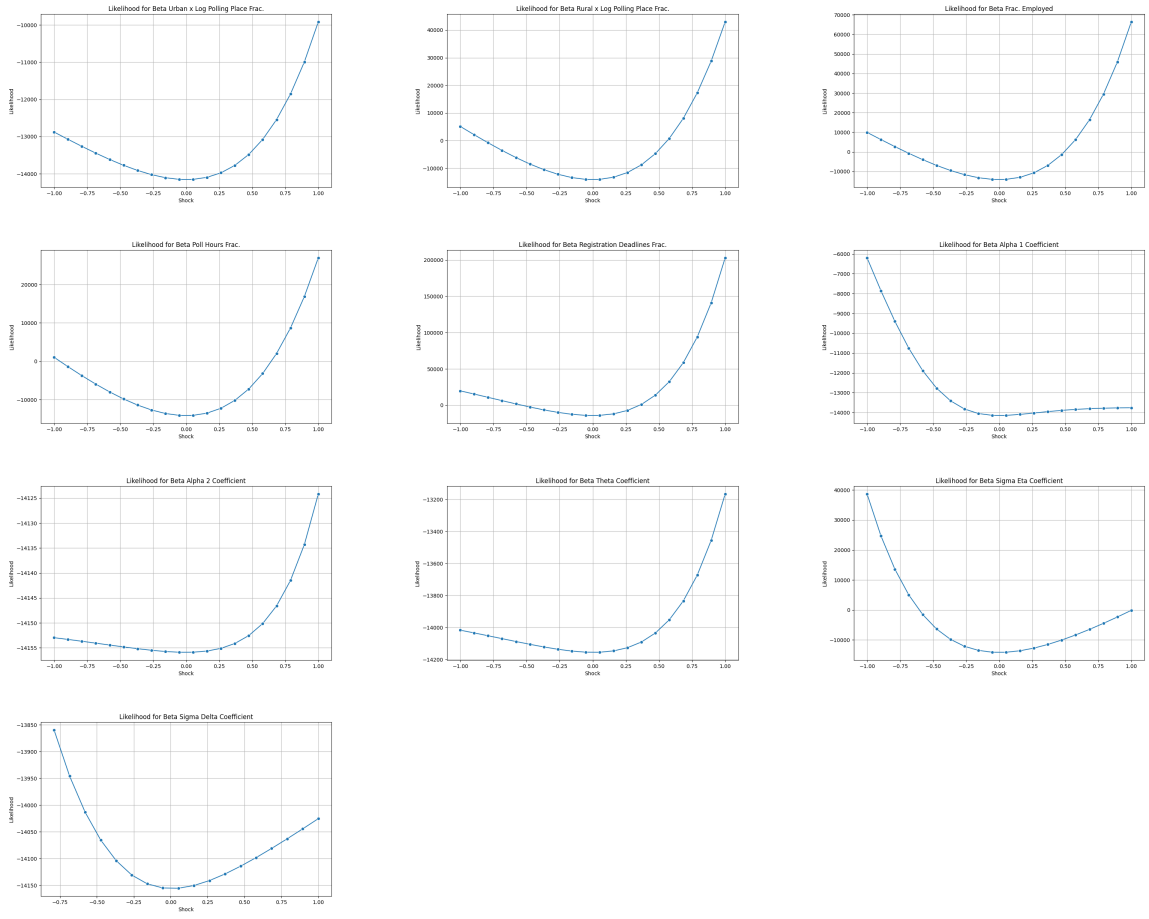

Log-likelihood evaluated around each β_k by applying additive shocks. Each panel holds all other coefficients fixed and perturbs only the k-th element.

$$F_{j_s,D}(\mathbf{e},\boldsymbol{\sigma}) \equiv \sigma_{j_s,D} - H\left(m(e_{s,D},e_{s,R}) - \mu_{j_s} - \frac{c_{j_s}}{p(\kappa_s)} + \eta_{j_s} + \delta_s - \zeta_{j_s}\right) = 0, \tag{22}$$

$$F_{j_s,R}(\mathbf{e},\boldsymbol{\sigma}) \equiv \sigma_{j_s,R} - H\left(m(e_{s,R},e_{s,D}) + \mu_{j_s} - \frac{c_{j_s}}{p(\kappa_s)} - \eta_{j_s} - \delta_s - \zeta_{j_s}\right) = 0, \tag{23}$$

where

$$\sigma_{s,p} = \sum_{j_s \in J_s} w_{j_s}\sigma_{j_s,p}, \qquad p(\kappa_s) := p(\sigma_{s,D},\sigma_{s,R}).$$

Stacking equations (22) and (23) over all counties in a given state yields the vector-valued function $F(\mathbf{e},\boldsymbol{\sigma}) \in \boldsymbol{R}^{2|J|}$. The Jacobians $\partial F/\partial\boldsymbol{\sigma}$ and $\partial F/\partial\mathbf{e}$ enter the implicit function theorem:

$$\frac{\partial\boldsymbol{\sigma}}{\partial\mathbf{e}} = -\left(\frac{\partial F}{\partial\boldsymbol{\sigma}}\right)^{-1}\left(\frac{\partial F}{\partial\mathbf{e}}\right),$$

which is used to compute marginal turnout responses to changes in campaign effort.

From these derivatives, I compute county-level elasticities of the form

$$\varepsilon_{j_s,p,q} = \left(\frac{\partial\sigma_{j_s,p}}{\partial e_{s,q}}\right)\left(\frac{e_{s,q}}{\sigma_{j_s,p}}\right), \qquad p,q \in \{D,R\},$$

G Validating the Effects of Competitiveness

G.1 Balance Tests and Covariate Adjustment

Before estimating the main regression, I test whether counties on opposite sides of a state border that share a media market differ systematically in observable characteristics. For each of fourteen covariates X_{cpt}, I estimate a regression of the form:

$$X_{cpt} = \beta \cdot \text{HighComp}_{cpt} + \delta_{pt} + \varepsilon_{cpt} \tag{24}$$

The indicator HighComp_{cpt} equals one if county c lies on the more competitive side of its border pair p in year t. All specifications include border-pair year fixed effects δ_{pt}. Standard errors are clustered at the county level.

Across the covariates tested, which include demographics, educational attainment, income, and employment status, only one shows statistically significant differences at conventional levels, the share of Hispanic residents is higher in competitive counties (0.6 percentage points, $p < 0.01$). While statistically significant, this difference is small in magnitude. As a robustness check, I separately run the main regression including this covariate, along with the full set of covariates used in the balance tests.

Results using pre-election polling averages instead of realized vote shares similarly demonstrate no significant differences in covariates across border counties, with the exception of fraction of residents employed.

Table 29: Covariate Balance Across Border Counties

	Realized Vote Shares		**Pre-Election Polls**	
Covariate	Estimate	Std. Error	Estimate	Std. Error
Male	0.001	(0.001)	0.001	(0.001)
Age 18–29	−0.001	(0.001)	−0.001	(0.001)
Age 65+	0.001	(0.001)	0.002	(0.001)
White	0.002	(0.005)	0.006	(0.004)
Black	−0.003	(0.003)	−0.004	(0.003)
American Indian/Alaska Native	−0.003	(0.003)	−0.004	(0.003)
Asian	0.001	(0.001)	0.001	(0.001)
Hispanic	0.006**	(0.002)	0.003	(0.002)
High School Only	0.001	(0.002)	0.004	(0.002)
Some College	0.002	(0.002)	0.001	(0.002)
College Only	0.001	(0.002)	0.000	(0.002)
College+	−0.001	(0.001)	−0.001	(0.001)
Log Median Income	0.012	(0.008)	0.011	(0.007)
Employed	0.004	(0.002)	0.004*	(0.002)
Observations	8,126		8,032	

Notes: Each row reports the coefficient from a separate regression of the specified covariate on an indicator for high competitiveness. All regressions include border-pair-by-year fixed effects. Covariates represent shares unless otherwise noted. Standard errors (in parentheses) are clustered by county. Significance levels: * p<0.05, ** p<0.01, *** p<0.001.

G.2 Robustness to Field Offices and Events

A potential concern with the border discontinuity design is that shared media markets may not fully equalize campaign exposure. In particular, ground operations such as field offices, canvassing hubs, or campaign events may vary discontinuously at state lines. To evaluate this concern, I construct a county-level dataset of field office and event activity from the 2008 to 2020 presidential elections, using disbursement records from the Federal Election Commission (FEC), available at `https://www.fec.gov/data/browse-data/?tab=bulk-data`. The data include transaction-level operating expenditures by presidential candidates' authorized committees.

I restrict attention to disbursements classified as rent, lease, or event-related, and exclude entries referring to equipment, services, or transportation using a set of keyword-based filters (for example, "car rental" or "audio/video"). Each transaction is mapped to a county using a ZIP-to-county crosswalk. If a ZIP code spans multiple counties, I conservatively assign the spending to all relevant counties.

The final dataset defines a binary indicator for whether any field office or event activity occurred in a given county-year. A balance test analogous to Appendix G.1 shows that more competitive counties are approximately five percentage points more likely to exhibit such activity, a difference that is statistically significant at the $p < 0.001$ level.

To assess the impact of this potential confound, I re-estimate the main regression after excluding any county pair where either county recorded ground activity. This removes approximately 1,000 county-border pair-year observations. As shown in Table 30, the estimated effect of competitiveness on turnout remains highly stable, with coefficients ranging from 0.053 to 0.067 depending on the specification. These results suggest that the main estimates are not driven by differences in ground operations across state lines.

Table 30: Effect of Competitiveness on Turnout (No-Office Sample Only)

		(1)	(2)	(3)
Competitiveness	Estimate	0.053***	0.057***	0.067***
	(Std. Error)	(0.011)	(0.011)	(0.010)
Controls		No	Balance-Test Sig.	All
Border Pair by Year FE		Yes	Yes	Yes
State FE		Yes	Yes	Yes
Observations		7,056	7,056	7,056
R^2_{within}		0.003	0.152	0.531

Notes: Each column reports regression estimates of the effect of state-level competitiveness on turnout, restricted to counties with no observed field offices or campaign events. All models include border-pair-by-year fixed effects and state fixed effects. Columns (2) and (3) sequentially add controls for covariates flagged as imbalanced in the balance tests and the full set of demographic and economic covariates. Standard errors (in parentheses) are clustered at the county-pair level. Significance levels: * p<0.05, ** p<0.01, *** p<0.001.

G.3 Digital Campaign Spending and Turnout

To further validate the model's predictions, I examine the relationship between state-level competitiveness, digital campaign spending, and turnout. I construct a state-level dataset of digital advertising expenditures for the 2020 election using records from the Center for Responsive Politics (`OpenSecrets.org`). I focus on the four largest general-election committees: Trump Make America Great Again Committee, Donald J. Trump for President, Biden for President, and the Biden Victory Fund. The first two are Republican committees, while the latter two are Democratic. I extract state-level totals of digital advertising from their public dashboards, merge the totals to obtain party-level spending, and compute per-capita values by dividing state totals by the voting-age population.

I then re-estimate the border discontinuity design described in equation 18, including per-capita digital spending as an additional regressor. The results, reported in Table 31, show that competitiveness remains a positive and statistically significant determinant of turnout even after controlling for digital spending.

Because digital advertising was minimal in earlier elections, I replicate the analysis using data from the 2008 and 2012 presidential contests. This limits the possibility that the estimated effect of competitiveness is driven by variation in digital spending. The results, reported in Table 32, yield coefficients nearly identical to the baseline estimates.

Table 31: Effect of Competitiveness and Digital Spending on Turnout

		(1)	(2)	(3)	(4)
Competitiveness	Estimate	0.153***	0.141***	0.171***	0.133***
	(Std. Error)	(0.013)	(0.019)	(0.018)	(0.014)
Digital Spending (per capita)	Estimate		0.006	−0.002	0.003
	(Std. Error)		(0.006)	(0.005)	(0.004)
Controls		No	No	Balance-Test Sig.	All
Border Pair by Year FE		Yes	Yes	Yes	Yes
State FE		No	No	No	No
Observations		2,044	2,044	2,044	2,044
R^2_{within}		0.111	0.112	0.229	0.581

Notes: Each column reports regression estimates of the effect of state-level competitiveness on turnout, with and without controls for per-capita digital advertising spending. All specifications include border-pair-by-year and state fixed effects. Columns (3) and (4) add covariates flagged as imbalanced in the balance tests and the full set of demographic and economic controls, respectively. Standard errors (in parentheses) are clustered at the county-pair level. Significance levels: * p<0.05, ** p<0.01, *** p<0.001.

Table 32: Effect of Competitiveness on Turnout (2008–2012 Elections)

		(1)	(2)	(3)
Competitiveness	Estimate	0.058***	0.061***	0.063***
	(Std. Error)	(0.009)	(0.009)	(0.010)
Controls		No Controls	Balance Sig. Controls	All Controls
Border Pair by Year FE		Yes	Yes	Yes
State FE		Yes	Yes	Yes
Observations		4,088	4,088	4,088
R^2_{within}		0.002	0.144	0.536

Notes: Each column reports regression estimates of the effect of state-level competitiveness on turnout in the 2008 and 2012 elections. All specifications include border-pair-by-year and state fixed effects. Column (2) adds Hispanic population share, the only covariate flagged as imbalanced in balance tests. Column (3) includes the full set of demographic and economic controls. Standard errors (in parentheses) are clustered at the county-pair level. Significance levels: * p<0.05, ** p<0.01, *** p<0.001.

G.4 Validation with Pre-Election Polling

To match with the model as closely as possible, the main regression uses realized vote shares to measure competitiveness. However, this approach may be subject to endogeneity concerns, as the same factors that drive turnout may also influence vote shares. To address this, I conduct a robustness check using pre-election polling data to measure competitiveness. I obtain state-level pre-election polling data from the *Fivethirtyeight* GitHub repository, which compiles polling averages from various sources and adjusts them for pollster quality, sample type, and recency. They give a predicted two-party vote share for each state in each election year over the election cycle throughout the 2008–2020 period. Polling data are unavailable for Delaware, Mississippi, and Wyoming in 2012. I use the final pre-election polling average available before Election Day for each state-year.

Similarly to the main regression, I define competitiveness as the ratio of the expected Democratic vote share to the expected Republican vote share in each state-year. The results

using pre-election polling data are reported in Table 33. While the overall effect is smaller, the estimated effect of competitiveness on turnout remains positive and statistically significant, with coefficients ranging from 0.027 to 0.030 depending on the specification. This translates to an increase in turnout of approximately 0.88 to 0.98 percentage points when moving from the non-battleground state mean ($\kappa_{st} = 0.677$) to full competitiveness ($\kappa_{st} = 1$), as shown in Table 34.

Although these estimates are lower than those from the main specification (1.72 to 2.17 percentage points) and the model's predicted effect (1.81 points), they remain directionally consistent and statistically robust. The smaller magnitudes may reflect greater noise in polling-based measures of competitiveness, which are based on expectations rather than realized outcomes.

Table 33: Effect of Competitiveness on Turnout (Pre-Election Polls)

		(1)	(2)	(3)
Competitiveness	Estimate	0.029**	0.027**	0.030***
	(Std. Error)	(0.010)	(0.010)	(0.009)
Controls		No	Balance-Test Sig.	All
Border Pair by Year FE		Yes	Yes	Yes
State FE		Yes	Yes	Yes
Observations		8,032	8,032	8,032
R^2_{within}		0.001	0.007	0.540

Notes: Each column reports regression estimates of the effect of state-level competitiveness (measured using pre-election polling averages) on turnout, at the county-border pair-year level. All models include border-pair-by-year fixed effects and state fixed effects. Columns (2) and (3) sequentially add controls for covariates flagged as imbalanced in the balance tests and the full set of demographic and economic covariates. Standard errors (in parentheses) are clustered at the county-pair level. Significance levels: * p<0.05, ** p<0.01, *** p<0.001.

Table 34: Estimated Turnout Effect of Moving from Average Non-Battleground Competitiveness to Full Competitiveness (Pre-Election Polls)

	No Controls	Balance Sig. Controls	All Controls
Effect (pp)	0.93	0.88	0.98

Notes: Each entry reports the estimated increase in turnout (in percentage points) associated with raising competitiveness from the non-battleground state mean ($\kappa_{st} = 0.677$) to full competitiveness ($\kappa_{st} = 1$), based on the pre-election polling estimates in Table 33.

G.5 Replicating Spenkuch and Toniatti (2018)

I replicate the design of Spenkuch and Toniatti (2018), which compares counties on opposite sides of media market boundaries but within the same state. This holds competitiveness fixed while allowing campaign exposure to vary.

Consistent with their findings, I find that per-capita campaign spending has no effect on turnout. However, when I regress the difference in per-capita spending between Democratic and Republican campaigns on the corresponding difference in vote shares, the estimated effect is large, positive, and statistically significant.

This closely mirrors the core result of Spenkuch and Toniatti (2018), who find that campaign advertising persuades but does not mobilize. The fact that I recover similar estimates using their design suggests that the strong turnout effects in my main analysis reflect differences in research design rather than differences in data.

H Marginal Cost Derivation

Total votes for party q in state s are given by:

$$V_{s,q} = \sum_{j_s \in s} \text{VAP}_{j_s} \cdot \sigma_{j_s,q},$$

Table 35: Reduced-Form Effects of Campaign Spending

Panel A: Turnout

		(1)	(2)	(3)
Total Spending	Estimate	0.001	0.000	0.000
	(Std. Error)	(0.001)	(0.001)	(0.000)
Controls		No	Balance-Test Sig.	All
Border Pair by Year FE		Yes	Yes	Yes
Observations		17,332	17,332	17,332
R^2_{within}		0.000	0.080	0.594

Panel B: Vote Share Difference

		(1)	(2)	(3)
Spending Difference	Estimate	0.018^{***}	0.015^{***}	0.010^{***}
	(Std. Error)	(0.005)	(0.004)	(0.003)
Controls		No	Balance-Test Sig.	All
Border Pair by Year FE		Yes	Yes	Yes
Observations		17,332	17,332	17,332
R^2_{within}		0.004	0.181	0.645

Notes: Panel A regresses county-level turnout on total per-capita campaign spending. Panel B regresses the difference in Democratic and Republican vote shares on the difference in per-capita campaign spending. All models include either border-pair or border-pair-by-year fixed effects. Columns (2) include only covariates flagged as imbalanced in balance tests. Columns (3) include the full set of demographic and economic controls. Standard errors (in parentheses) are clustered by county. Significance levels: * p<0.05, ** p<0.01, *** p<0.001.

where VAP_{j_s} is the voting-age population in county j_s, and $\sigma_{j_s,q}$ is the party-specific turnout share. Differentiating yields:

$$\frac{\partial V_{s,q}}{\partial e_{s,q}} = \sum_{j_s \in s} \text{VAP}_{j_s} \cdot \frac{\partial \sigma_{j_s,q}}{\partial e_{s,q}},$$

and thus:

$$MCV_{s,q} = \left(\sum_{j_s \in s} \text{VAP}_{j_s} \cdot \frac{\partial \sigma_{j_s,q}}{\partial e_{s,q}} \right)^{-1}.$$

References

Aggarwal, Minali et al. (2023). "A 2 million-person, campaign-wide field experiment shows how digital advertising affects voter turnout". In: *Nature Human Behaviour* 7.3, pp. 332–341.

Ainsworth, Robert, Emanuel Garcia Munoz, and Andres Munoz Gomez (2022). "District competitiveness increases voter turnout: evidence from repeated redistricting in North Carolina". In: *University of Florida*.

Bagwe, Gaurav, Juan Margitic, and Allison Stashko (2022). *Polling Place Location and the Costs of Voting*. Tech. rep. Technical Report. https://jmargitic. github. io/JM/Margitic_JMP. pdf.

Bär, Domink, Nicolas Pröllochs, and Stefan Feuerriegel (2025). "The role of social media ads for election outcomes: Evidence from the 2021 German election". In: *PNAS Nexus*, pgaf073.

Bouton, Laurent et al. (2023). *Pack-crack-pack: gerrymandering with differential turnout*. Tech. rep. National Bureau of Economic Research.

Burden, Barry C et al. (2014). "Election laws, mobilization, and turnout: The unanticipated consequences of election reform". In: *American Journal of Political Science* 58.1, pp. 95–109.

Bursztyn, Leonardo et al. (2024). "Identifying the effect of election closeness on voter turnout: Evidence from swiss referenda". In: *Journal of the European Economic Association* 22.2, pp. 876–914.

Cantoni, Enrico and Vincent Pons (2021). "Strict ID laws don't stop voters: Evidence from a US nationwide panel, 2008–2018". In: *The Quarterly Journal of Economics* 136.4, pp. 2615–2660.

Coate, Stephen and Michael Conlin (2004). "A group rule—utilitarian approach to voter turnout: Theory and evidence". In: *American Economic Review* 94.5, pp. 1476–1504.

Coate, Stephen, Michael Conlin, and Andrea Moro (2008). "The performance of pivotal-voter models in small-scale elections: Evidence from Texas liquor referenda". In: *Journal of Public Economics* 92.3-4, pp. 582–596.

Degan, Arianna and Antonio Merlo (2011). "A structural model of turnout and voting in multiple elections". In: *Journal of the European Economic Association* 9.2, pp. 209–245.

Dixit, Avinash and John Londregan (1996). "The determinants of success of special interests in redistributive politics". In: *the Journal of Politics* 58.4, pp. 1132–1155.

Downs, Anthony (1957). "An economic theory of political action in a democracy". In: *Journal of Political Economy* 65.2, pp. 135–150.

Duffy, John and Margit Tavits (2008). "Beliefs and voting decisions: A test of the pivotal voter model". In: *American Journal of Political Science* 52.3, pp. 603–618.

Enos, Ryan D and Anthony Fowler (2014). "Pivotality and turnout: Evidence from a field experiment in the aftermath of a tied election". In: *Political Science Research and Methods* 2.2, pp. 309–319.

— (2018). "Aggregate effects of large-scale campaigns on voter turnout". In: *Political Science Research and Methods* 6.4, pp. 733–751.

Fraga, Bernard L, Daniel J Moskowitz, and Benjamin Schneer (2022). "Partisan alignment increases voter turnout: Evidence from redistricting". In: *Political behavior* 44.4, pp. 1883–1910.

Gerber, Alan S (2004). "Does campaign spending work? Field experiments provide evidence and suggest new theory". In: *American Behavioral Scientist* 47.5, pp. 541–574.

Gerber, Alan S et al. (2011). "How large and long-lasting are the persuasive effects of televised campaign ads? Results from a randomized field experiment". In: *American Political Science Review* 105.1, pp. 135–150.

Grumbach, Jacob M and Charlotte Hill (2022). "Rock the registration: Same day registration increases turnout of young voters". In: *The Journal of Politics* 84.1, pp. 405–417.

Herrera, Helios, David K Levine, and Cesar Martinelli (2008). "Policy platforms, campaign spending and voter participation". In: *Journal of Public Economics* 92.3-4, pp. 501–513.

Kalla, Joshua L and David E Broockman (2018). "The minimal persuasive effects of campaign contact in general elections: Evidence from 49 field experiments". In: *American Political Science Review* 112.1, pp. 148–166.

Kawai, Kei, Yuta Toyama, and Yasutora Watanabe (2021). "Voter turnout and preference aggregation". In: *American Economic Journal: Microeconomics* 13.4, pp. 548–586.

Levine, David K and Thomas R Palfrey (2007). "The paradox of voter participation? A laboratory study". In: *American political science Review* 101.1, pp. 143–158.

Li, Quan, Michael J Pomante, and Scot Schraufnagel (2018). "Cost of voting in the American states". In: *Election Law Journal: Rules, Politics, and Policy* 17.3, pp. 234–247.

Lindbeck, Assar and Jörgen W Weibull (1987). "Balanced-budget redistribution as the outcome of political competition". In: *Public choice* 52.3, pp. 273–297.

Nickerson, David W (2006). "Volunteer phone calls can increase turnout: Evidence from eight field experiments". In: *American Politics Research* 34.3, pp. 271–292.

Palfrey, Thomas R and Howard Rosenthal (1983). "A strategic calculus of voting". In: *Public choice* 41.1, pp. 7–53.

Schuster, Steven Sprick (2020). "Does campaign spending affect election outcomes? New evidence from transaction-level disbursement data". In: *The Journal of Politics* 82.4, pp. 1502–1515.

Shachar, Ron and Barry Nalebuff (1999). "Follow the leader: Theory and evidence on political participation". In: *American Economic Review* 89.3, pp. 525–547.

Sides, John, Lynn Vavreck, and Christopher Warshaw (2022). "The effect of television advertising in United States elections". In: *American Political Science Review* 116.2, pp. 702–718.

Spenkuch, Jörg L and David Toniatti (2018). "Political advertising and election results". In: *The Quarterly Journal of Economics* 133.4, pp. 1981–2036.

Strömberg, David (2008). "How the Electoral College influences campaigns and policy: the probability of being Florida". In: *American Economic Review* 98.3, pp. 769–807.