

Motivated or Mobilized?

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

Luke Miller

Georgetown University

November 9, 2025

(Online appendix)

For latest version see [here](#)

For online appendix see [here](#)

A Miscellaneous Tables and Figures

Table 1: 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
	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

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 2: Estimated Coefficients: Race and Ethnicity

Parameter	Variable	Coefficient	Std. Error	P-Value
μ	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
	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

Notes: See Table 1.

Table 3: 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
	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 1.

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:

$$\mathbb{E}[U(\text{abstain})] = \Pr(R \text{ wins} | i \text{ abstains})u_{iR} + \Pr(D \text{ wins} | i \text{ abstains})u_{iD}. \quad (1)$$

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

$$\mathbb{E}[U(\text{vote } R)] = \Pr(R \text{ wins} | i \text{ votes } R)u_{iR} + \Pr(D \text{ wins} | i \text{ votes } R)u_{iD} - c_i. \quad (2)$$

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

$$\mathbb{E}[U(\text{vote } R)] - \mathbb{E}[U(\text{abstain})] > 0. \quad (3)$$

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

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

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

Substituting equations (1) and (2) into (3) gives the following condition for participation:

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

Rearranging terms by candidate utility:

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

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.

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

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

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 (5) and (6)), as presented in the main text.

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

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

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} :

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

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:

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

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:


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

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 ($\hat{\pi}_s$) 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 1 compares the histograms of these simulated outcomes against the normal distribution with the implied mean and variance. 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.




(a) 2008 Election



(b) 2012 Election



(c) 2016 Election



(d) 2020 Election

Figure 1: 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.

B.4 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:

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

and compute two probabilities:

$$\begin{aligned}\pi_D^{\text{full}} &= \Pr[\sigma_{sD}(\eta, \zeta, \delta) > \sigma_{sR}(\eta, \zeta, \delta)], \\ \pi_D^{\text{approx}} &= \Pr[\sigma_{sD}(0, 0, \delta) > \sigma_{sR}(0, 0, \delta)],\end{aligned}$$

where $\sigma_{sD}(\cdot)$ and $\sigma_{sR}(\cdot)$ denote the state-level turnout functions for Democrats and Republicans, respectively, defined in Equations (5) and (6). 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:

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

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

This section verifies that the threshold value $\hat{\delta}_s$ defined in 10 is unique for a broad range of parameter values and campaign effort profiles.¹

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 $N_{\text{sim}} = 1,000$

¹In principle, two distinct fixed points may exist: an interior solution and a degenerate boundary solution of the form $(x, 0, 0)$, where x satisfies 10 under $\sigma_{s,D} = \sigma_{s,R} = 0$. I exclude such degenerate solutions from the analysis, as they do not correspond to meaningful equilibria.

Monte Carlo simulations, I attempt to solve for $(\hat{\delta}_s, \sigma_{s,D}, \sigma_{s,R})$ from 10, 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:

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

where $\text{Var}_t(\cdot)$ 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 4.

Table 4: Dispersion of $(\hat{\delta}_s, \sigma_{s,D}, \sigma_{s,R})$ across Random Initializations

	$\hat{\delta}_s$	$\sigma_{s,D}$	$\sigma_{s,R}$
$\max_r \text{Var}_t(\cdot)$	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 10 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.

B.6 State-Level Win Probabilities and Their Derivatives

For each effort profile (e_{sD}, e_{sR}) , the state-level win probabilities $\tilde{\pi}_s(e_{sD}, e_{sR})$ are determined implicitly from the voter-turnout equilibrium. Specifically, each $\tilde{\pi}_s$ depends on the value of $\hat{\delta}_{s,t}$ that equalizes expected Democratic and Republican vote shares within state s :

$$\sigma_{j_s D} = \sigma_{j_s R}, \quad \text{where} \quad \sigma_{j_s D} = \sum_{j_s} w_{j_s} H \left(m(e_{sD}, e_{sR}) - \mu_{j_s} - \frac{c_{j_s}}{p(\sigma_{j_s D}, \sigma_{j_s R})} + \hat{\delta}_s \right),$$

and $\sigma_{j_s R}$ is defined analogously. Thus, to calculate the FOC and the Jacobian, we need the jacobian and hessian of $\tilde{\pi}_s(e_{sD}, e_{sR})$ with respect to effort. Turnout rates are themselves calculated via Newton's method applied to the voter-turnout equilibrium within each state:

$$\begin{aligned} F(\hat{\delta}, \sigma_{j_s D}, \sigma_{j_s R}) &= \begin{bmatrix} F_1(\hat{\delta}, \sigma_{j_s D}, \sigma_{j_s R}) \\ F_2(\hat{\delta}, \sigma_{j_s D}, \sigma_{j_s R}) \\ F_3(\hat{\delta}, \sigma_{j_s D}, \sigma_{j_s R}) \end{bmatrix} \\ &= \begin{bmatrix} \sum_{j_s} w_{j_s} H \left(v_{j_s}^D - \frac{c_{j_s}}{p(\sigma_{j_s D}, \sigma_{j_s R})} + \hat{\delta} \right) - \sum_{j_s} w_{j_s} H \left(v_{j_s}^R - \frac{c_{j_s}}{p(\sigma_{j_s D}, \sigma_{j_s R})} - \hat{\delta} \right) \\ \sum_{j_s} w_{j_s} H \left(v_{j_s}^D - \frac{c_{j_s}}{p(\sigma_{j_s D}, \sigma_{j_s R})} + \hat{\delta} \right) - \sigma_{j_s D} \\ \sum_{j_s} w_{j_s} H \left(v_{j_s}^R - \frac{c_{j_s}}{p(\sigma_{j_s D}, \sigma_{j_s R})} - \hat{\delta} \right) - \sigma_{j_s R} \end{bmatrix} = \mathbf{0}. \end{aligned} \tag{4}$$

where j_s indexes counties with weights w_{j_s} in state s , H is a CDF with density $h = H'$, and

$$v_{j_s}^D = m(e_{sD}, e_{sR}) - \mu_{j_s}, \quad v_{j_s}^R = m(e_{sD}, e_{sR}) + \mu_{j_s}.$$

I use $p(\sigma_{j_s D}, \sigma_{j_s R})$ instead of $p(\kappa_s)$ here to emphasize that the turnout rates $\sigma_{j_s D}$ and $\sigma_{j_s R}$ are the objects that enter into the voting efficacy function.

The first equation, $F_1 = 0$, ensures that the tie shock $\hat{\delta}$ equalizes expected vote shares, while $F_2 = 0$ and $F_3 = 0$ ensure that the turnout rates $\sigma_{j_s D}$ and $\sigma_{j_s R}$ that enter into the

voting efficacy function $p(\sigma_{jsD}, \sigma_{jsR})$ are consistent with county-level turnout decisions.

To compute the gradient of $\tilde{\pi}_s$ efficiently, I apply the implicit function theorem (IFT) to the system defining the voter-turnout equilibrium, avoiding the need for automatic differentiation through nested fixed-point iterations. Within the Jacobian for the candidate first-order conditions (Eq. 10), I also require the derivatives of these gradients with respect to effort, which are obtained by second-order implicit differentiation.

Differentiating this system with respect to the candidates' efforts and applying the implicit function theorem gives the total derivative:

$$\frac{d(\hat{\delta}, \sigma_{jsD}, \sigma_{jsR})}{d(e_{sD}, e_{sR})} = - \left(\frac{\partial F}{\partial(\hat{\delta}, \sigma_{jsD}, \sigma_{jsR})} \right)^{-1} \frac{\partial F}{\partial(e_{sD}, e_{sR})}.$$

I evaluate this expression numerically by solving the full three-equation system at each iteration when calculating optimal effort levels. This approach ensures that both the tie condition and the turnout identities are satisfied when computing gradients of the state-level win probabilities.

The probability of winning state s is given by

$$\tilde{\pi}_s(e_{sD}, e_{sR}) = 1 - G(\hat{\delta}_s) \tag{5}$$

where G is the CDF of the state-level shock δ_s . In other words, for a given $\hat{\delta}_s$, if the state-level shock δ_s is greater than $\hat{\delta}_s$, then the Democratic candidate wins the state. Combining this expression with the total derivative of $(\hat{\delta}, \sigma_{jsD}, \sigma_{jsR})$ with respect to (e_{sD}, e_{sR}) gives the gradients of the state-level win probabilities

$$\tilde{\pi}_s^1 = -g(\hat{\delta}_s) \frac{d\hat{\delta}_s}{de_{sD}}, \quad \tilde{\pi}_s^2 = -g(\hat{\delta}_s) \frac{d\hat{\delta}_s}{de_{sR}},$$

where g is the density corresponding to G , and $d\hat{\delta}_s/de_{sD}$ and $d\hat{\delta}_s/de_{sR}$ are obtained from the total derivative above.

Likewise, the second derivatives of the state-level win probabilities with respect to effort are obtained by differentiating the gradients:

$$\begin{aligned}\tilde{\pi}_s^{11} &= -g'(\hat{\delta}_s) \left(\frac{d\hat{\delta}_s}{de_{sD}} \right)^2 - g(\hat{\delta}_s) \frac{d^2\hat{\delta}_s}{de_{sD}^2}, \\ \tilde{\pi}_s^{22} &= -g'(\hat{\delta}_s) \left(\frac{d\hat{\delta}_s}{de_{sR}} \right)^2 - g(\hat{\delta}_s) \frac{d^2\hat{\delta}_s}{de_{sR}^2}, \\ \tilde{\pi}_s^{12} &= -g'(\hat{\delta}_s) \frac{d\hat{\delta}_s}{de_{sD}} \frac{d\hat{\delta}_s}{de_{sR}} - g(\hat{\delta}_s) \frac{d^2\hat{\delta}_s}{de_{sD}de_{sR}}.\end{aligned}$$

The second derivatives of $\hat{\delta}_s$ with respect to effort are obtained by using the second-order IFT on the system in (4).

B.7 Solving for the Equilibrium Effort Allocations

The equilibrium effort allocations (e_{sD}^*, e_{sR}^*) are obtained by jointly solving the first-order conditions implied by each candidate's optimization problem, subject to their respective budget constraints. This ensures that the model's predicted effort levels are mutually optimal given expectations about state-level competitiveness and voter responsiveness.

In the following, I deviate slightly from the main text by letting e_{sD} and e_{sR} denote total effort in state s (rather than effort per capita); $\tilde{\pi}_{s(esD, e_{sR})}$ converts these to per-capita units internally. This is purely a notational change. Let BG denote the set of battleground states and $S \equiv |BG|$.

Candidate D chooses effort allocations $\{e_{sD}\}_{s \in BG}$ to maximize the probability of securing at least 270 electoral votes:

$$\max_{\{e_{sD}\}_{s \in BG}} \Phi \left(\frac{\sum_{s \in BG} \tilde{\pi}_s(e_{sD}, e_{sR}) l_s - (270 - EV_D)}{\sqrt{\sum_{s \in BG} \tilde{\pi}_s(e_{sD}, e_{sR}) [1 - \tilde{\pi}_s(e_{sD}, e_{sR})] l_s^2}} \right) \text{ s.t. } \sum_{s \in BG} e_{sD} \leq B_D. \quad (6)$$

Likewise, for the Republican candidate:

$$\max_{\{e_{sR}\}_{s \in BG}} \Phi \left(\frac{\sum_{s \in BG} [1 - \tilde{\pi}_s(e_{sD}, e_{sR})] l_s - (270 - EV_R)}{\sqrt{\sum_{s \in BG} \tilde{\pi}_s(e_{sD}, e_{sR}) [1 - \tilde{\pi}_s(e_{sD}, e_{sR})] l_s^2}} \right) \text{ s.t. } \sum_{s \in BG} e_{sR} \leq B_R. \quad (7)$$

These inequalities bind in equilibrium, so I solve the equivalent equality-constrained problem. Let $\tilde{\pi}_s^1 \equiv \partial \tilde{\pi}_s / \partial e_{sD}$ and $\tilde{\pi}_s^2 \equiv \partial \tilde{\pi}_s / \partial e_{sR}$, and for compactness, define

$$f^D(\mathbf{e}_D, \mathbf{e}_R) = \frac{\sum_{s \in BG} \tilde{\pi}_s l_s - (270 - EV_D)}{\sqrt{\sum_{s \in BG} \tilde{\pi}_s (1 - \tilde{\pi}_s) l_s^2}}, \quad f^R(\mathbf{e}_D, \mathbf{e}_R) = \frac{\sum_{s \in BG} [1 - \tilde{\pi}_s] l_s - (270 - EV_R)}{\sqrt{\sum_{s \in BG} \tilde{\pi}_s (1 - \tilde{\pi}_s) l_s^2}},$$

Then the first-order condition for candidate D in state s is

$$\phi(f^D(\mathbf{e}_D, \mathbf{e}_R)) \left[\frac{l_s \tilde{\pi}_s^1}{\sqrt{\sum_{t \in BG} \tilde{\pi}_t (1 - \tilde{\pi}_t) l_t^2}} - \frac{\sum_{t \in BG} \tilde{\pi}_t l_t - (270 - EV_D)}{2 [\sum_{t \in BG} \tilde{\pi}_t (1 - \tilde{\pi}_t) l_t^2]^{3/2}} l_s^2 \tilde{\pi}_s^1 (1 - 2\tilde{\pi}_s) \right] = \lambda_D, \quad (8)$$

where λ_D is the Lagrange multiplier on the budget constraint and $\phi(\cdot)$ is the standard normal density. An analogous condition holds for candidate R :

$$\phi(f^R(\mathbf{e}_D, \mathbf{e}_R)) \left[\frac{-l_s \tilde{\pi}_s^2}{\sqrt{\sum_{t \in BG} \tilde{\pi}_t (1 - \tilde{\pi}_t) l_t^2}} - \frac{\sum_{t \in BG} (1 - \tilde{\pi}_t) l_t - (270 - EV_R)}{2 [\sum_{t \in BG} \tilde{\pi}_t (1 - \tilde{\pi}_t) l_t^2]^{3/2}} l_s^2 \tilde{\pi}_s^2 (1 - 2\tilde{\pi}_s) \right] = \lambda_R. \quad (9)$$

The system in (8)–(9) is solved using a hybrid root-finder: a Newton–Raphson step is first attempted; if the residual norm does not contract or yields negative effort levels, the algorithm switches to a damped Newton (Levenberg–Marquardt) update in log-space, where $x_{s,q} = \log e_{s,q}$ ensures $e_{s,q} \geq 0$ by construction.

The first-order conditions can be written succinctly as

$$\phi(f^D(\mathbf{e}_D, \mathbf{e}_R)) \frac{\partial f^D}{\partial e_{sD}} = \lambda_D, \quad \phi(f^R(\mathbf{e}_D, \mathbf{e}_R)) \frac{\partial f^R}{\partial e_{sR}} = \lambda_R.$$

Newton’s method and Levenberg–Marquardt attempt to find the root of the stacked

system:

$$\begin{bmatrix} \phi(f^D) \frac{\partial f^D}{\partial e_{1D}} - \lambda_D \\ \vdots \\ \phi(f^D) \frac{\partial f^D}{\partial e_{sD}} - \lambda_D \\ \phi(f^R) \frac{\partial f^R}{\partial e_{1R}} - \lambda_R \\ \vdots \\ \phi(f^R) \frac{\partial f^R}{\partial e_{sR}} - \lambda_R \\ \sum_{s \in BG} e_{sD} - B_D \\ \sum_{s \in BG} e_{sR} - B_R \end{bmatrix} = \mathbf{0}.$$

The first S equations correspond to the Democratic candidate's FOCs, and the next S to the Republican candidate's, followed by two budget equalities.

To simplify the problem, I subtract each element of the candidate's FOC system by the next one, yielding $|BG| - 1$ equations per candidate that do not depend on the Lagrange multipliers. The last two equations are the budget constraints, which do depend on the multipliers. This also allows me to divide out the $\phi(f^D)$ and $\phi(f^R)$ from each row. This transformation improves numerical stability and speeds convergence. The system is now:

$$\begin{bmatrix} \frac{\partial f^D}{\partial e_{1D}} - \frac{\partial f^D}{\partial e_{2D}} \\ \vdots \\ \frac{\partial f^D}{\partial e_{(s-1)D}} - \frac{\partial f^D}{\partial e_{sD}} \\ \frac{\partial f^R}{\partial e_{1R}} - \frac{\partial f^R}{\partial e_{2R}} \\ \vdots \\ \frac{\partial f^R}{\partial e_{(s-1)R}} - \frac{\partial f^R}{\partial e_{sR}} \\ \sum_{s \in BG} e_{sD} - B_D \\ \sum_{s \in BG} e_{sR} - B_R \end{bmatrix} = \mathbf{0}.$$

This yields a square system of dimension $2S$, with one equation per unknown effort level.

Let,

$$g_s^D = \frac{\partial f^D}{\partial e_{sD}} - \frac{\partial f^D}{\partial e_{(s+1)D}}, \quad g_s^R = \frac{\partial f^R}{\partial e_{sR}} - \frac{\partial f^R}{\partial e_{(s+1)R}}, \quad h_D = \sum_{s \in BG} e_{sD} - B_D, \quad h_R = \sum_{s \in BG} e_{sR} - B_R.$$

The Jacobian of the $2S$ equations with respect to the $2S$ unknowns is

$$J(\mathbf{e}_D, \mathbf{e}_R) = \left[\begin{array}{ccc|ccc} \frac{\partial g_1^D}{\partial e_{1D}} & \dots & \frac{\partial g_1^D}{\partial e_{SD}} & \frac{\partial g_1^D}{\partial e_{1R}} & \dots & \frac{\partial g_1^D}{\partial e_{SR}} \\ \vdots & & \vdots & \vdots & & \vdots \\ \frac{\partial g_{S-1}^D}{\partial e_{1D}} & \dots & \frac{\partial g_{S-1}^D}{\partial e_{SD}} & \frac{\partial g_{S-1}^D}{\partial e_{1R}} & \dots & \frac{\partial g_{S-1}^D}{\partial e_{SR}} \\ \frac{\partial g_1^R}{\partial e_{1D}} & \dots & \frac{\partial g_1^R}{\partial e_{SD}} & \frac{\partial g_1^R}{\partial e_{1R}} & \dots & \frac{\partial g_1^R}{\partial e_{SR}} \\ \vdots & & \vdots & \vdots & & \vdots \\ \frac{\partial g_{S-1}^R}{\partial e_{1D}} & \dots & \frac{\partial g_{S-1}^R}{\partial e_{SD}} & \frac{\partial g_{S-1}^R}{\partial e_{1R}} & \dots & \frac{\partial g_{S-1}^R}{\partial e_{SR}} \\ \hline 1 & \dots & 1 & 0 & \dots & 0 \\ 0 & \dots & 0 & 1 & \dots & 1 \end{array} \right], \quad (10)$$

which is a $2S \times 2S$ matrix. In the next section (B.6), I describe how to compute the derivatives of the state-level win probabilities that enter into these expressions.

C Identification Appendix

C.1 Identification via Monte Carlo Simulation

For the first validation exercise, I use the observed county- and state-level covariates from the empirical application, (X_{js}^μ, X_{js}^c) , 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 $(e_{s,D}^{(r)}, e_{s,R}^{(r)})$, 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 5 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 5: 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.

C.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 $k = 1, \dots, K$, 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 2 - 5 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.

D Data Appendix

D.1 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 ,

Figure 2: 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 3: 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 4: 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 5: 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.

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 impute 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 $R^2 = 0.838$.

D.2 Summary statistics

Table 6: Summary Statistics for Covariates Used in the Partisan-Bias Parameter (Battle-ground 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 7: 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 8: Summary Statistics for Cost of Voting Index by Urban/Rural Status and State Type

Statistic	Urban Counties		Rural Counties	
	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 9: Summary Statistics for Election-Law Indices by State Type

Statistic	Voter ID Index		Poll Hours Index		Reg. Deadline Index	
	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 10: Summary Statistics for Cost of Voting Covariates by State Type


Statistic	Fraction Employed	
	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.

D.3 Battleground state classification

To illustrate the sharp concentration of campaign resources, Figure 6 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 6: 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 11 and 12 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 11: 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 12: 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.4 Constructing Campaign Budget Shares and Total Effort

D.5 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 $TV_{s,p}$ 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.6 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” → “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.7 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 $\text{media} \succ \text{online} \succ \text{print} \succ \text{telemarketing} \succ \text{travel}$, ensuring each transaction appears exactly once.

E Results Appendix

E.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 (5) and (6). In equilibrium, the following system must be satisfied for each county j_s in state s and each party $p \in \{D, R\}$:

$$F_{j_s,D}(\mathbf{e}, \boldsymbol{\sigma}) \equiv \sigma_{j_s,D} - H\left(m(e_{s,D}, e_{s,R}) - \mu_{j_s} - \frac{c_{j_s}}{p(\kappa_s)} + \eta_{j_s} + \delta_s - \zeta_{j_s}\right) = 0, \quad (11)$$

$$F_{j_s,R}(\mathbf{e}, \boldsymbol{\sigma}) \equiv \sigma_{j_s,R} - H\left(m(e_{s,R}, e_{s,D}) + \mu_{j_s} - \frac{c_{j_s}}{p(\kappa_s)} - \eta_{j_s} - \delta_s - \zeta_{j_s}\right) = 0, \quad (12)$$

where

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

Stacking equations (11) and (12) over all counties in a given state yields the vector-valued function $F(\mathbf{e}, \boldsymbol{\sigma}) \in \mathbf{R}^{2|J|}$. 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

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

F Validating the Effects of Competitiveness using a Border Discontinuity Design

F.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:

$$X_{cpt} = \beta \cdot \text{HighComp}_{cpt} + \delta_{pt} + \varepsilon_{cpt} \quad (13)$$

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 13: Covariate Balance Across Border Counties

Covariate	Realized Vote Shares		Pre-Election Polls	
	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: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

F.2 Results

Table 14 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.

I additionally test whether results are sensitive to the measurement of competitiveness. Instead of realized vote shares, Appendix F.5 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 mea-

asures of closeness.

Table 14: Effect of Competitiveness on Turnout

		(1)	(2)	(3)
Competitiveness	Estimate (Std. Error)	0.054*** (0.007)	0.059*** (0.007)	0.064*** (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 15: 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 14.

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 F.6 for details.

F.3 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 F.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 16, 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 16: 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: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

F.4 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](https://www.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 19, including per-capita digital spending as an additional regressor. The results, reported in Table 17, 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 18, yield coefficients nearly identical to the baseline estimates.

Table 17: 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 18: 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: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

F.5 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 19. 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 20.

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 19: 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 20: 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 19.

F.6 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.

G 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 21: Reduced-Form Effects of Campaign Spending

Panel A: Turnout

		(1)	(2)	(3)
Total Spending	Estimate (Std. Error)	0.001 (0.001)	0.000 (0.001)	0.000 (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 (Std. Error)	0.018*** (0.005)	0.015*** (0.004)	0.010*** (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:

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

H Computational Strategy for the Popular Vote

The primary source of computational complexity in the national popular vote simulation is the cross-state dependency created by the national efficacy term $p(\sigma_D, \sigma_R)$. In the baseline model, state-level win probabilities could be calculated independently conditional on effort. This meant I could approximate the total number of electoral votes using a normal distribution. Now, however, a shock δ_s in one state affects turnout in all other states by altering the national vote totals and, consequently, the perceived efficacy of voting.

To calculate the optimal campaign strategies under this new scenario, I employ a Monte Carlo simulation approach to approximate the probability of winning the national popular vote. For a given effort allocation (e_D, e_R) , I draw M vectors of state-level shocks $\{\delta_s\}_{s \in \mathcal{S}}$. Given there are 50 states, I draw 50 state-level shocks for each $m \in M$ vector.

I calculate the optimal campaign strategies in two steps. First, I use a gradient-based optimizer to find an approximate equilibrium. Then, I refine this solution using an iterative best-response (IBR) procedure. In both cases, I use a gradient-based optimization method (the Adam optimizer) to find the optimal effort allocations for each candidate. However, the Monte Carlo procedure yields a discontinuous and non-differentiable objective function, preventing the use of standard gradient-based optimization methods. To address this, I approximate the indicator function for winning with a logistic sigmoid function, creating a “soft win” objective:

$$W_D(e_D, e_R) \approx \frac{1}{M} \sum_{m=1}^M \frac{1}{1 + \exp(-k \cdot \text{margin}^{(m)})}$$

where $\text{margin}^{(m)} = \sigma_D^{(m)} - \sigma_R^{(m)}$ is the national popular vote share margin in simulation draw m , and $\sigma_p^{(m)}$ is the total national vote share for party p given the m -th vector of state-level shocks. This objective function is smooth, allowing for the computation of exact gradients via automatic differentiation.

In the first step, I use the Adam optimizer to perform simultaneous gradient updates

to find an approximate equilibrium. This phase incorporates an annealing schedule for the sigmoid’s steepness parameter, k , which is incrementally increased from $k = 10$ to $k = 25$, and finally to $k = 50$. In practice, this means the objective function starts as a smooth approximation of the win probability and becomes increasingly sharp, approaching the true indicator function as k increases. During this phase, I calculate the gradient estimates using a mini-batch approach, using only 50 simulation draws per gradient evaluation, while using $M = 1000$ draws to evaluate the objective function itself. This technique balances computational efficiency with the need for accurate gradient estimates.


In the second step, I refine the initial solution using an iterative best-response (IBR) procedure with k fixed at 50 and use all M draws to calculate the gradient. However, given the higher number of draws used for the gradient, I limit $M = 500$. In each step of the IBR procedure, one candidate’s strategy is held fixed while the other’s optimal response is found using the Adam optimizer. To enhance stability, the IBR updates are damped using a mixing parameter, such that the strategy for the next iteration is a weighted average of the current strategy and the newly computed best response.

Figure 7 visualizes the convergence of this algorithm for each party in each election cycle. Each plot shows the infinity norm of the difference between a campaign’s effort allocation vector across successive iterations, with the y-axis on a logarithmic scale to better visualize the approach to zero. The left panel of each plot shows the Warmup Phase using the simultaneous gradient-based optimizer with annealing, while the right panel displays the subsequent Iterative Best Response (IBR) Phase. For each election year (2008, 2012, 2016, and 2020), there are two plots: one for the Democratic candidate (left) and one for the Republican candidate (right).

As the plots demonstrate, the norm difference consistently trends downward across all scenarios, and by the final iterations it typically falls below 10^{-4} . For comparison, the normalized total budget constraint is between 20-30. This indicates that the campaign effort allocations are stabilizing, suggesting convergence to an equilibrium strategy profile under

the national popular vote system.

Figure 7: Convergence of Campaign Effort under a National Popular Vote System



Notes: The figure visualizes the convergence of the computational algorithm used to find the equilibrium effort allocations under the national popular vote system. Each plot displays the infinity norm of the difference between a campaign's effort allocation vector across successive iterations, with the y-axis on a logarithmic scale. The eight panels show the results for each party (Democratic, Republican) in each election cycle (2008, 2012, 2016, 2020). Each plot is split into two phases: “Phase 1: Warmup” uses a simultaneous gradient-based optimizer with an annealing schedule, while “Phase 2: IBR” refines the solution using an iterative best-response (IBR) procedure. The consistent downward trend in the norm difference, which typically falls below 10^{-4} by the final iteration, indicates that the campaign effort allocations are stabilizing and have converged to an equilibrium strategy profile.