These observations follow from equation (2) and known results in the literature on zero-sum games. In particular, because on-path spending levels do not depend on past realizations of popularity, the candidates' equilibrium spending paths would be the same even if popularity was not fully observable. Furthermore, because the game is zero-sum, the equilibrium path of play is unique and robust to allowing candidates to move sequentially within a period, with arbitrary order of moves (see, for example, Mertens et al., 2015).

2.3 Equilibrium Spending Ratios

To say more about the equilibrium spending paths and the candidates' equilibrium probabilities of winning, we need to impose additional assumptions on how spending levels affect the popularity process. Under the following assumption, we can fully specify the equilibrium evolution of the popularity process.

Assumption 2. *The function p is homogeneous of degree $\beta \geq 0$.*

The function $p(x, y) = \alpha_1 x^\beta - \alpha_2 y^\beta$ satisfies this assumption and further satisfies Assumption 1 when $\beta \in (0, 1)$ and $\alpha_1, \alpha_2 > 0$. Another example that satisfies the assumption is the function $p(x, y) = \alpha(\log x - \log y)$, which is homogeneous of degree 0 and satisfies Assumption 1 when $\alpha > 0$.[10]

We define the *spending ratio* of a candidate in period t to be the ratio between her spending level in period t and the remaining budget available to her that period: in period t, candidate 1's spending ratio is x_t/X_t and candidate 2's is y_t/Y_t. We refer to the ratio of spending in period $t+1$ to spending in period t for a candidate as the *consecutive period spending ratio*, and we use $r_{1,t} := x_{t+1}/x_t$ and $r_{2,t} := y_{t+1}/y_t$ from here on to denote them.

Assumption 2 implies two key results that inform our analysis of spending patterns in actual elections. The first is an equal spending ratio result. It says that the candidates' spending ratios equal each other on the path of play. The second is a constant spending growth result. It says that the candidates' consecutive period spending ratios equal the same constant in all periods.

Proposition 2. *Suppose Assumptions 1 and 2 hold. Then, in the unique equilibrium path of the dynamic campaigning game,*

[10]For this example, the model is not closed since p is not defined when either $x = 0$ or $y = 0$. However, we can close the model by assuming that: (i) if any candidate i spends 0 at any time t, then the game ends immediately with candidate $j \neq i$ winning so long as j spends a positive amount, and (ii) if both candidates simultaneously spend 0, then the game ends with each candidate winning with probability 1/2. The results of Proposition 1 and 2 continue to hold under this amendment.

(i) the candidates' spending ratios equal each other every period: $x_t/X_t = y_t/Y_t$ *for all periods* t.

(ii) the candidates' consecutive period spending ratios equal each other and are constant through time; in particular, $r_{1,t} = r_{2,t} = \delta^{1/(\beta-1)}$ *for all periods* $t < T-1$.

The results in Proposition 2 are based on the following reasoning. To maximize the probability of winning the election, both candidates equalize the (decay-weighted) marginal benefit of spending at any period $t < T-1$ with the (decay-weighted) marginal benefit of spending in period $T-1$, just ahead of the election. The homogeneity of function p (Assumption 2) implies that the ratio of candidates' first order conditions depends only on the ratio of their spending levels at time t (x_t/y_t) and at time $T-1$ (x_{T-1}/y_{T-1}). The equal spending ratio result then follows iteratively from the strict concavity of p and from the budget balance condition, which implies

$$\frac{x_{t-1}}{y_{t-1}} = \frac{X_{t-1}}{Y_{t-1}} = \frac{X_t - \sum_{t\leq t'<T-1} x_{t'}}{Y_t - \sum_{t\leq t'<T-1} y_{t'}}.$$

The equal spending ratio result implies some additional equilibrium properties. For example, it implies that the candidates' on-path consecutive period spending ratios are equal to each other, i.e. $r_{1,t} = r_{2,t}$ for all periods t. In addition, because budgets are fixed, the equal spending ratio result implies that the ratio x_t/y_t of the candidates' spending levels in any period t (which we refer to as the *cross-candidate spending ratio*) is a constant that is equal to the ratio of the starting budgets; that is, $x_t/y_t = X_0/Y_0$ for all periods t. The two properties described above yield part (ii) of Proposition 2.

Given that δ does not exceed 1, Proposition 2 implies that candidates' spending grows over time if both $\beta, \delta < 1$.[11] The expression in the proposition verifies that if $\delta = 1$ (i.e., if popularity leads do not decay), then the candidates spread their budgets evenly across periods. Since p is concave, the candidates want to smooth their spending over time. The lack of decay further implies that this smoothing is full: candidates allocate the same share of their initial budgets to each period. On the other hand, if $\delta < 1$, then spending increases over time, and the fraction of the initial budget each candidate spends at time t is

$$\gamma_t = \frac{x_t}{X_0} = \frac{y_t}{Y_0} = \frac{r-1}{r^T - 1} r^t, \tag{4}$$

[11]Although we assume $\delta \leq 1$, none of the above results relies on this assumption. If $\delta > 1$, popularity leads tend to amplify over time; and, on the equilibrium path, the candidates would decrease their spending over time if $\beta < 1$.

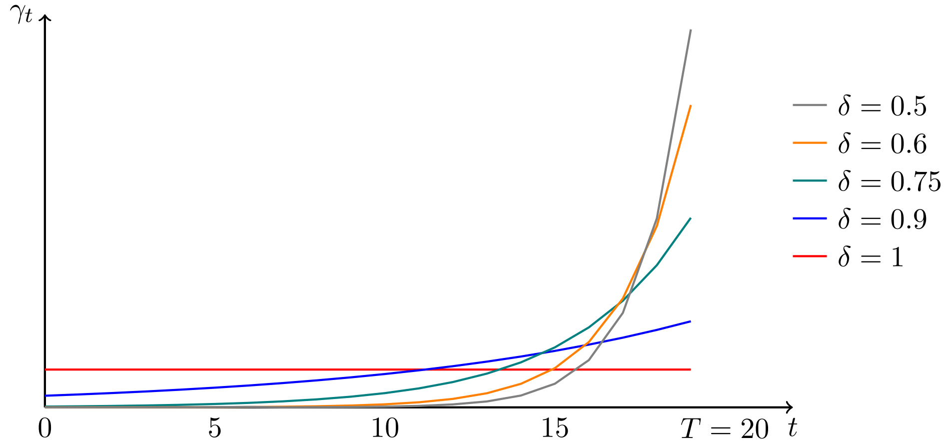

Figure 2: Equilibrium share of the initial budget $\gamma_t = x_t/X_0$ that the candidates spend over time when $T = 20$, $\beta = 0$ and δ takes different values.

where $r = \delta^{1/(\beta-1)}$ is the common consecutive period spending ratio. In this case, the decay of popularity leads generates a force that pushes for greater spending in later periods.

The comparative statics of γ_t reflect these countervailing forces. If β increases, the marginal return to spending decreases at a slower rate within each period. Candidates thus spend even more towards the end of the campaign and less in the early stages. As $\beta \to 1^-$ candidates spend all of their resources in the final period. As δ decreases, popularity leads decay more and candidates have an incentive to invest less in the early stages of the race and more in the later stages. Figure 2 depicts this last feature, plotting γ_t for $\beta = 0$ and different values of δ.

Strategic Spending Considerations. A candidate's optimal spending behavior varies with the spending behavior of the other candidate only if the effects of the spending levels on the drift of the popularity process (i.e., function p) are not additively separable. Suppose $p(x, y) = (x - y)/2(x + y)$, which is not additively separable.[12] Given any behavior by candidate 2, the first order condition for candidate 1 implies that the marginal benefit to spending in period $t < T - 1$ equals the marginal benefit to spending in the final period $T - 1$:

$$\delta^{T-1-t} \frac{y_t}{(x_t + y_t)^2} = \frac{y_{T-1}}{(x_{T-1} + y_{T-1})^2}.$$

[12]To close the model when both candidates spend 0, see footnote 10. In addition, although this function does not satisfy Assumption 1(c), the results of Propositions 1 and 2 hold with $\beta = 0$; in particular, the first order conditions are satisfied at an interior equilibrium, and since the function is homogeneous of degree 0 the common consecutive period spending ratio is $r = 1/\delta$.

Both sides of this equation feature expressions of the form $y/(x+y)^2$, whose partial derivative in y is $(x-y)/(x+y)^3$ and in x is $-2y/(x+y)^3$. With this in mind, suppose that candidate 2 marginally lowers his spending in period t and, to keep his budget balanced, increases his spending in a later period, say $T-1$. The previous observation implies that candidate 1's best response at time t would be to either increase her spending (this happens if $x_t \leq y_t$), or to lower it as well but by a factor smaller than candidate 2's (this happens if $x_t > y_t$).[13] (From this we see that a candidate's optimal response to a change in the opponent's spending depends not just on the direction of the change, but also on its magnitude.) In both cases, the cross-candidate spending ratio x_t/y_t increases. Analogously, if candidate 2 raises her spending in any period t relative to the equilibrium level, and candidate 1 best responds, then the cross-candidate spending ratio x_t/y_t decreases.

Suppose that candidate 2 naively spends almost all of her budget in the final period. The observations above imply that a strategic candidate 1 would best respond by spending a positive amount in all periods and increasing her spending over time at a rate that is faster than the equilibrium rate, i.e. the rate stated in Proposition 2 for $\beta = 0$. On the other hand, if candidate 2 naively allocates his budget evenly across all periods, a strategic candidate 1 would best respond by increasing her spending over time at a slower rate than the equilibrium rate.

2.4 Discussion

Our baseline model provides a useful benchmark to understand how strategic candidates compete against each other in a dynamic setting. To highlight the key forces behind this dynamic contest, it abstracts away from several factors that shape spending in actual elections. Yet, our theoretical framework is flexible enough to accommodate several of these factors. For instance, advantages (or disadvantages) due to incumbency, to prior legislative records, or to a candidate's name recognition can affect the initial lead in relative popularity, z_0, or starting budgets, X_0 and Y_0.

Candidates can also differ in the effectiveness of their campaign spending. These differences may depend on differences in how their campaigns are organized, or on the fact that one candidate is simply better than the other at campaigning. A candidate's policy

[13] To see why, note that if candidate 2 lowers his spending in period t from y_t to αy_t with $\alpha < 1$ and candidate 1 also lowers her spending from x_t to αx_t (or to an even lower amount) then $y_t/(x_t + y_t)^2$ drops to $y_t/[\alpha(x_t + y_t)^2]$ (or even lower) and the FOC is violated. For the FOC to hold, candidate 1's spending level at t must be larger than αx_t.

platform may also be more popular than the one of the other candidate. We can capture these features through asymmetries in the partial first derivatives of p.

Although the payoffs we have assumed imply a winner-take-all electoral rule, our equilibrium analysis immediately extends to the case where the candidates' payoffs are linear (or piecewise linear) in relative popularity on election day, z_T. Therefore, it covers the case in which the margin of victory also matters to the candidates.

Equation (1) assumes that relative popularity evolves according to an AR(1) process. In the Online Appendix, we examine non-separable popularity processes, imposing additional assumptions to guarantee that the first-order approach is still sufficient to characterize the equilibrium evolution of relative popularity.

Our baseline model assumes that candidates have fixed budgets, or equivalently that they can forecast exactly how much money they will have by the end of the campaign, and they are not allowed to finish in debt. In the Online Appendix, we consider a variant of the model in which budgets are uncertain and evolve over time in response to fluctuations in relative popularity. We show for a specification of that model that the equal spending ratio result continues to hold but the constant spending growth result does not. Because spending decisions depend on the candidates' expectations of how their budgets evolve and because these expectations vary with fluctuations in relative popularity, equilibrium spending also evolves stochastically.

Finally, in the Online Appendix we present a model in which we allow voters to react to campaign spending differently, following the approach of the marketing literature. Our model of the electorate gives rise to a popularity process for the candidates that is equivalent to equation (1). We then demonstrate how this approach can be used to derive policy implications; specifically, we study the welfare effects of campaign silence laws and of spending caps.

In the next section, we look at three additional variants of our model.

3 Variants

3.1 Early Voting

In the baseline model, the candidates' payoffs depend only on their relative popularity on election day, i.e., at time T. But in many elections voters can and do cast their votes prior to election day, which suggests that the candidates' payoffs should depend on realizations of relative popularity even prior to time T. We now analyze how early voting affects the candidates' spending decisions.

Consider the baseline model, but now suppose that voters can vote from period $\hat{T} < T$ onwards. Suppose that the difference in votes cast for the two candidates in each period $t \geq \hat{T}$ is proportional to their relative popularity in that period, Z_t, and let the number of votes cast in period $t \geq \hat{T}$ be a proportion $\xi \in (0, 1)$ of the number of votes cast in period $t + 1$. Turnout thus increases over time at a constant rate. This assumption simplifies the notation, but our analysis extends to other assumptions concerning the evolution of turnout so long as candidates can perfectly forecast turnout rates and cannot manipulate them. As ξ converges to zero, almost all votes are cast at time T and we converge to the baseline model. Finally, we assume that despite early voting, either candidate is still able to eventually win the election if she is sufficiently more popular than her opponent at date T, no matter how low her popularity was in previous periods.[14]

In each period t, candidate 1 thus maximizes $\Pr[\sum_{t=\hat{T}}^{T} \xi^{T-t} Z_t \geq 0 \mid (z_{t'}, X_{t'}, Y_{t'})_{t' \leq t}]$, while candidate 2 minimizes this expression. An analogue to problem (3) in the baseline model holds in this variant as well. In particular, given candidate 2's on path spending levels $\{y_0, ..., y_{T-1}\}$, candidate 1's equilibrium spending path $\{x_0, ..., x_{T-1}\}$ now maximizes

$$\sum_{t=0}^{\hat{T}-1} \sum_{t'=0}^{T-\hat{T}} \xi^{t'} \delta^{T-1-t-t'} p(x_t, y_t) + \sum_{t=\hat{T}}^{T-1} \sum_{t'=0}^{T-1-t} \xi^{t'} \delta^{T-1-t-t'} p(x_t, y_t), \tag{5}$$

subject to the same nonnegativity and budget constraints as in problem (3). Candidate 2's spending path correspondingly minimizes this expression subject to her own nonnegativity and budget constraints.

Proposition 3. *Suppose Assumptions 1 and 2 hold. Then in the unique equilibrium path of the game with early voting,*

(i) $x_t/X_t = y_t/Y_t$ *for all periods* t.

(ii) the consecutive period spending ratios equal each other: $r_{1,t} = r_{2,t} = r_t$ *for all periods* t*, and in particular,* $r_t = \delta^{1/(\beta-1)}$ *if* $t < \hat{T} - 1$*, and*

$$r_t = \left[\delta\left(1 + \frac{1}{\sum_{t'=0}^{T-2-t} \xi^{-(T-1-t-t')} \delta^{T-2-t-t'}}\right)\right]^{1/(\beta-1)} \quad \textit{if } t \geq \hat{T} - 1$$

[14] This condition holds if $\xi(2 - \xi^{T-\hat{T}}) < 1$, which is implied by $\xi < 1/2$. Alternatively, we could also assume that candidate 1 maximizes (and candidate 2 minimizes) the difference in candidate 1 and 2's vote shares, $\sum_{t=\hat{T}}^{T} \xi^{T-t} Z_t$. The results of Proposition 3 extend to this case.

Proposition 3 asserts that under early voting candidates continue to allocate the same share of their budgets on the path of play. But early voting modifies the spending path. Because the term inside the large round bracket in Proposition 3 is larger than 1 (and because $\beta < 1$), the spending path is flatter than in the baseline model ($r_t < r$ when $t \geq \hat{T} - 1$). As some voters vote early, candidates now have a new incentive to allocate a larger share of their budget to earlier periods, relative to the baseline model. Moreover, once early voting begins, the consecutive period spending ratio is decreasing in ξ. As the share of voters who vote in early periods increases (higher ξ), the candidates' spending levels will be more evenly distributed (lower r_t for $t > \hat{T} - 1$).

3.2 Valuing Money Left Over

In the variants studied so far, the two candidates are purely office-motivated and fully deplete their budgets by the end of the race because they do not value money left over. In reality, money left over may be valuable: candidates may want to save for future campaigns, or for investments outside politics—to the extent that this is legally allowed.

To capture this, let X_T and Y_T be money left over at the end of the campaign for candidates 1 and 2 respectively. Assume that in each period t candidate 1 maximizes $\Pr[Z_T \geq 0 \mid (z_{t'}, X_{t'}, Y_{t'})_{t' \leq t}] + \kappa_1 X_T$, while candidate 2 maximizes $(1 - \Pr[Z_T \geq 0 \mid (z_{t'}, X_{t'}, Y_{t'})_{t' \leq t}]) + \kappa_2 Y_T$. The parameter $\kappa_i > 0$ reflects candidate i's marginal value for money. On top of saving money and benefiting from this at rate κ_i, we also assume that each candidate i can overspend his budget by borrowing money at a cost equal to κ_i.[15] Thus, X_T and Y_T can be negative. In addition, for tractability we assume that Assumption 2 holds with $\beta = 0$ and we define the function q so that $p(x, y) = p(x/y, 1) =: q(x/y)$ for $y > 0$. (To close the model in the case of $y = 0$, see footnote 10.)

In this variant of the model, candidates trade off spending on the campaign against not spending on it. The marginal benefit of spending depends on the probability of winning, which is history-dependent as it varies with the popularity shocks. The marginal value of not spending on the race is, on the other hand, history-independent. The marginal rate of substitution between spending in a given period and not spending on the campaign is thus history-dependent. As a result, the candidates' equilibrium paths of spending depend on the realization of the popularity shocks. In the baseline model, by contrast, candidates have no incentive to not spend their money on the race. Popularity shocks thus affect the

[15] To simplify the analysis, we abstract from the time dimension when we model borrowing: a unit of money borrowed at any point during the race has the same cost κ_i.

marginal benefit of spending money in all periods by the same amount and the marginal rate of substitution across periods is independent of these shocks.

When money left over is valuable, spending by both candidates decreases as the race becomes more lopsided. To state this popularity dependence formally, define the following quantity for every time t:

$$\zeta((\varepsilon_{t'})_{t'=0}^{t-1}) = \frac{\sum_{t'=0}^{T-1} \delta^{T-1-t'} q\,(\kappa_1/\kappa_2) + \delta^T z_0 + \sum_{t'=0}^{t-1} \delta^{T-1-t'} \varepsilon_{t'}}{\sigma \sqrt{\sum_{t'=t}^{T-1} \delta^{T-1-t'}}} \tag{6}$$

This quantity measures the expected electoral advantage that one candidate has over the other at time t: when one candidate has a large popularity advantage over the other, $|\zeta((\varepsilon_{t'})_{t'=0}^{t-1})|$ is large.

Proposition 4. *Suppose Assumptions 1 and 2 hold with $\beta = 0$. Then in the unique equilibrium path of the game in which candidates' marginal valuations for money left over are $\kappa_1, \kappa_2 > 0$,*

(i) $x_t/y_t = \kappa_2/\kappa_1$ for all periods t,

(ii) x_t and y_t are both decreasing in $|\zeta((\varepsilon_{t'})_{t'=0}^{t-1})|$, and

(iii) if κ_1 and κ_2 decrease by the same factor for both candidates, then x_t and y_t increase for both candidate in all periods t and for all realizations z_t.

Part (i) of Proposition 4 says that all the equilibrium cross-candidate spending ratios x_t/y_t equal the ratio of marginal valuations of money left over, κ_2/κ_1. (Recall that in the baseline model, all the cross-candidate spending ratios equal the ratio of starting budgets X_0/Y_0.) In equilibrium, both candidates equalize the marginal benefit of spending with its opportunity cost, which is now equal to the candidates' marginal value of money left over. When function p is homogeneous of degree 0, the ratio of the candidates' marginal benefits to spending in any given period depends only on the cross-candidate spending ratio. In equilibrium, the cross-candidate spending ratio, x_t/y_t, must then equal the ratio of marginal values of money left over.

Part (ii) of the proposition says that spending by both candidates decreases as the election becomes more lopsided, implying that the candidates' spending levels are no longer independent of relative popularity. In particular, if candidate 1 becomes more popular relative to candidate 2, then candidate 2 prefers to save more of her budget because her probability of winning is now smaller. In equilibrium, this pushes candidate 1 to lower her spending as well. This finding is in line with the "discouragement effect" studied in the dynamic contest literature (see Konrad, 2009 and Fu and Wu, 2019 for reviews). In our setting, the result arises from the existence of an outside option (i.e., saving money for after the campaign) that becomes more appealing for a candidate as her odds of winning worsen.

Finally, part (iii) of the proposition asserts that when the marginal values of money κ_1 and κ_2 decrease proportionally for the two candidates, the spending levels of the two candidates go up uniformly in each period. This implies that high stakes elections (those with lower κ_1 and κ_2) should see on average higher spending.

3.3 Targetable Subpopulations

In any campaign, candidates choose not just when to spend their resources, but also how to target these resources across voters—for example by targeting specific geographic areas or media markets. Suppose that the two candidates compete over a set of targetable subpopulations. The set of subpopulations is $\{1, 2, ..., S\}$ and the payoffs of the candidates depend on how these different subpopulations aggregate.

Popularity in each subpopulation s is represented by the random variable Z_t^s with realizations z_t^s. We assume that $(Z_t^s)_s$ are distributed according to a multivariate normal distribution with arbitrary variance-covariance matrix. For each subpopulation s, the popularity process is

$$Z_{t+1}^s = p(x_t^s, y_t^s) + \delta^s Z_t^s + \varepsilon_t^s, \tag{7}$$

where $\varepsilon_t^s \sim \mathcal{N}(0, (\sigma^s)^2)$ and these shocks are iid over time. Each subpopulation s thus has its own decay parameter δ^s, and its own variance $(\sigma^s)^2$. In addition, as in the previous section, we assume that the function p satisfies Assumptions 1 and 2 with $\beta = 0$, so that $p(x, y) = p(x/y, 1) = q(x/y)$ for some function q.[16]

The aggregation rule for the outcomes in the various subpopulations is arbitrary, but we impose the following assumptions: the candidates' payoffs depend only on the vector $(Z_T^s)_{s=1}^S$, the game is still zero sum, and candidate 1's payoff is strictly increasing in each Z_T^s, while candidate 2's is strictly decreasing in each Z_T^s. More formally, denote candidate 1's payoff $u((Z_T^s)_{s=1}^S)$ so that candidate 2's payoff is $-u((Z_T^s)_{s=1}^S)$, and assume that

$$\frac{\partial u((Z_T^s)_{s=1}^S)}{\partial Z_T^s} > 0, \qquad \text{for every } s. \tag{8}$$

[16]We extend the assumption in footnote 10 as follows: if a candidate spends an amount equal to 0 in some subpopulation, then the game ends and the candidate wins with probability 1/2 if the other candidate is also spending an amount equal to 0 in some subpopulation, and loses with probability 1 otherwise.

For this model, we can show that the equal spending ratio result holds subpopulation by subpopulation, which is stated in part (i) of Proposition 5 below. However, unlike in the baseline model, spending decisions may depend on the history of the popularity processes. If the competition in some subpopulations becomes lopsided (in terms of the candidates' relative popularity), the marginal benefit of spending money in those subpopulations decreases for both candidates. Candidates will then react by concentrating their spending in other, more competitive subpopulations. Relative popularity within different populations thus plays a role in spending decisions.

This popularity-dependence does not arise in the special case in which payoffs are a weighted sum of relative popularity in each district at time T. In this case, candidate 1's marginal benefit of increasing her popularity in a specific subpopulation is constant and it is equal to the marginal benefit of candidate 2. Moreover, under this assumption, we can characterize the consecutive period spending ratios for this model as well as the optimal allocation of resources across districts in each period—results that are stated in parts (ii) and (iii) of Proposition 5 respectively. The following assumption, which strengthens the monotonicity assumption in equation (8), states the condition formally.

Assumption 3. *For weights $\{w^s\}_{s=1}^S$ such that $w^s > 0$ and $\sum_{s=1}^S w^s = 1$,*

$$u\left((Z_T^s)_{s=1}^S\right) = \sum_{s=1}^{S} w^s Z_T^s.$$

Assumption 3 fits either a setting where candidates allocate resources across multiple media markets, or one in which the candidates are two parties that compete to maximize the number of seats in a legislature, seats are allocated proportionally in each district, and the number of seats assigned to each district depends on the district population reflected in w^s.

To state Proposition 5, let h_t denote histories prior to the candidates choosing their period-t spending levels. Let the consecutive period spending ratios for the two candidates in any district s be $r_{1,t}^s = x_{t+1}^s / x_t^s$ and $r_{2,t}^s = y_{t+1}^s / y_t^s$.

Proposition 5. *Suppose Assumptions 1 and 2 hold with $\beta = 0$. In any equilibrium of this targetable subpopulations extension,*

(i) in each subpopulation s, $x_t^s / X_t = y_t^s / Y_t$ for every t.

(ii) if Assumption 3 holds, then, in each subpopulation, the candidates' consecutive period spending ratios conditional on any on-path history equal each other: $r^s_{1,t} = r^s_{2,t} = 1/\delta^s$ for all s and all histories h_t.

(iii) if Assumption 3 holds, then for all periods t and any pair of subpopulations s, s',

$$\frac{x^s_t}{x^{s'}_t} = \frac{y^s_t}{y^{s'}_t} = \frac{w^s}{w^{s'}} \left(\frac{\delta^s}{\delta^{s'}} \right)^{T-t-1}$$

By Proposition 5(iii), the allocation of resources across subpopulations given total spending in a period is independent of the popularity process. Moreover, candidates spend more on subpopulations that have greater electoral weight and for whom popularity leads decay at a slower rate. Finally, the differences in spending due to different decay rates are maximal at the beginning of the campaign and decrease as election day approaches. These results hold even if the candidates' investments in any one subpopulation also affect their relative popularity in other subpopulations.

4 TV Ad Spending in Actual Elections

We now look at actual campaign spending data through the lens of our baseline model. Under Assumptions 1 and 2, the predicted pattern of spending is given by $r_{1,t} = r_{2,t} = r = \delta^{1/(\beta-1)}$ (see Proposition 2 and Figure 2). Our main goal is to use this relationship to recover election-specific estimates of δ from patterns of spending. If candidates compete according to our baseline model, this gives us estimates of how they perceive the decay rate $1 - \delta$ when making their spending decisions. Actual spending paths obviously depend on some factors that our baseline model does not account for. Our estimation exercise, nevertheless, still informs us on how much we can explain with a simple competitive environment. As such, it can guide our understanding of what are the factors that are likely missing to get a better fit with the data.

Before proceeding, we introduce the data we use and we investigate the extent to which two important implications of our baseline model are violated in the data: the equal spending ratio result ($x_t/X_t = y_t/Y_t$ for all t) and the constant spending growth result ($r_{1,t} = r_{2,t} = r$ for all t).

4.1 Data

We focus on subnational American elections, namely U.S. House, Senate, and gubernatorial elections in the period 2000 to 2014.

Spending in our model refers to all spending—TV ads, calls, mailers, door-to-door canvasing visits—that directly affects the candidates' relative popularity. But for some of these categories, it is not straightforward to separate out the part of spending that has a direct impact on relative popularity from the part that does not (e.g. fixed administrative costs). For television ads, it is straightforward to do this, so we focus exclusively on TV ad spending. Television advertising constitutes around 35% of the total expenditures by congressional candidates, and around 90% of all advertisement expenditures during the period we study (see, e.g., Albert, 2017). Furthermore, for TV ads, we have access to the exact timing of the candidates' expenditures. We proceed under the assumption that any spending on other types of campaign activities that directly affect relative popularity is proportional to spending on TV ads.

Our TV ad spending data are from the Wesleyan Media Project and the Wisconsin Advertising Database, which draws information directly from TV channels. For each election in which TV ads were bought, the database contains information about the candidate that each ad supports, the date it was aired, and the estimated cost. The dataset does not include information on the source of spending (whether PACs or the candidates themselves), but the vast majority of expenditure on TV advertising is likely to happen through PACs (Martin and Peskowitz, 2018).

For the year 2000, the dataset covers only the 75 largest Designated Market Areas (DMAs), and for years 2002-2004 it covers the 100 largest DMAs. The data from 2006 onward covers all 210 DMAs. We obtain the amount spent on ads from total ads bought and price per ad. Ad price data are missing for 2006, so for that year we estimate prices using ad prices in 2008.[17]

We focus on races where the leading two candidates in terms of vote share are from the Democratic and the Republican party. We label the Democratic candidate as candidate 1 and the Republican candidate as candidate 2, so that x_t, X_0, etc. refer to the Democrat's spending, budget, etc. and y_t, Y_0, etc. refer to the Republican's.

[17] In principle, as election day approaches, TV ad prices can increase. Increases in total spending over time could confound price increases with increased advertising. Federal regulations, however, limit the ability of TV stations to increase ad prices close to elections. TV stations must charge political candidates "the lowest unit charge of the station for the same class and amount of time for the same period" (Chapter 5 of Title 47 of the United States Code 315, Subchapter III, Part 1, Section 315, 1934). These regulations allay some of these concerns.

In our model, spending decisions are made at discrete moments in time defined so that the inter-period decay rate $1 - \delta$ is constant. This raises the question of how to define a period of spending in the data, given that spending data are reported irregularly. To address this issue, in the Online Appendix we examine a continuous time formulation of our model in which candidates make spending decisions at fixed intervals of time and the decay rate is constant. There, we prove an identification result that implies that the level of aggregation of spending is irrelevant: e.g., if candidates make their spending decisions daily but the data are aggregated weekly, then the sum of what they spend over seven days is the same as in a setting in which they make spending decisions weekly. Given this result, we proceed by aggregating spending data at the weekly level.

To restrict attention to general elections, we focus on the 12 weeks leading to election day, though we drop the final week which is typically incomplete since elections are held on Tuesdays.[18] We exclude elections that are clearly not genuine contests to which our model does not apply. These are elections in which one of the candidates did not spend anything for more than half of the period studied. This leaves us with 346 House, 122 Senate, and 133 gubernatorial elections, tabulated in the Appendix. We define the total budgets of the candidates to be the total amount that they spend over these 12 weeks. In the Online Appendix, we replicate our analyses by excluding fewer elections (leaving us with 1163 elections over 14 years), or by allowing for a longer time window for each election (20 weeks instead of 12).

Table 1 reports summary statistics for the elections we consider. There is considerable difference in the amount spent between state-wide and House elections, with another key difference being the time at which candidates start spending positive amounts. For statewide races, candidates spend on average about \$6 million on TV ads, with most candidates already spending positive amounts 12 weeks prior to the election. For House races, they spend \$1.5 million on average and the majority of candidates start spending 9 weeks out.

In addition, there is variation in the amount spent by candidates competing in the same race. The average difference in the amount spent by the candidates competing in the same congressional election is one third of the average total spending for those races, while for gubernatorial elections the same difference exceeds half. Finally, candidates tend to spend more in more competitive elections: the overall amount spent is higher in elections where

[18]In some cases, primaries are held less than 12 weeks before the general election, but ad spending for the general election before the primaries is typically zero. In the rare cases where ad spending for primary elections happens, we exclude it from our analysis.

Table 1: Descriptive Statistics

	N	Open Seat Election	Incumbent Competing	No Excuse Early Voting	Average total spending	Average spending difference
Senate	122	68	54	82	6019 (5627)	1962 (2921)
Governor	133	59	74	92	5980 (9254)	3173 (6,337)
House	346	97	249	223	1533 (1304)	521 (615)
Overall	601	224	377	397	3428 (5581)	1401 (3,461)

Average Spending and Standard Deviations in Parentheses by Week and Election Type

Week	-11	-10	-9	-8	-7	-6	-5	-4	-3	-2	-1
Senate	196	250	266	314	357	477	545	652	716	860	1,002
	(291)	(328)	(403)	(487)	(401)	(505)	(577)	(724)	(803)	(947)	(1,047)
Share spending 0	0.270	0.180	0.123	0.082	0.008	0	0	0	0	0	0
Governor	262	253	258	316	420	416	530	597	701	800	1,019
	(632)	(468)	(424)	(581)	(865)	(579)	(1,249)	(1,015)	(1,305)	(1,523)	(1,956)
Share spending 0	0.297	0.207	0.139	0.068	0.030	0	0	0	0	0	0
House	17	27	38	56	83	120	137	177	212	250	303
	(41)	(55)	(57)	(85)	(93)	(134)	(134)	(182)	(219)	(270)	(340)
Share spending 0	0.653	0.545	0.386	0.246	0.095	0	0	0	0	0	0
Early Voting	113	123	128	168	223	262	320	390	449	526	624
	(324)	(256)	(256)	(348)	(488)	(404)	(694)	(663)	(775)	(895)	(1033)
No Early Voting	99	122	144	162	194	250	283	321	373	436	569
	(213)	(246)	(286)	(314)	(259)	(317)	(348)	(394)	(473)	(524)	(866)
Open Seat	164	183	191	217	279	324	362	445	521	602	729
	(404)	(311)	(325)	(352)	(476)	(445)	(485)	(635)	(800)	(892)	(1046)
Incumbent	75	87	99	135	174	219	275	320	366	432	532
	(189)	(202)	(218)	(324)	(386)	(324)	(657)	(550)	(606)	(716)	(931)
Close Election	122	131	154	200	250	292	383	479	544	661	858
	(318)	(236)	(320)	(421)	(539)	(407)	(915)	(791)	(746)	(884)	(1266)
Not Close Election	103	120	125	152	199	245	278	322	376	430	506
	(280)	(259)	(242)	(297)	(369)	(364)	(412)	(476)	(659)	(741)	(820)
Close Budgets	97	118	129	150	196	264	301	362	411	477	587
	(190)	(209)	(255)	(225)	(246)	(339)	(385)	(442)	(488)	(550)	(710)
Not Close Budgets	117	127	137	178	227	254	313	370	434	510	620
	(351)	(282)	(275)	(404)	(525)	(405)	(727)	(680)	(813)	(938)	(1149)

Note: Spending on television advertising for the twelve weeks prior to election dates, excluding the final (partial) week, as elections are held on Tuesdays. The upper panel reports the breakdown of elections that are open seat versus those that have an incumbent running, the number of elections in which voters can vote early without an excuse to do so, average spending levels by the candidates, and the average difference in spending between the two candidates, all by election type. The lower panel presents average spending for each week in our dataset, by election type. Standard deviations are in parentheses. All monetary amounts are in units of $1,000. Close elections are races where the final difference in vote shares between two candidates is less than 5 percentage points. Close budget races are those in which the ratio of budgets of the two candidates lies in the interval $(0.75, 1.25)$.

there is no incumbent, and in elections where the final margin of victory is thin. We will consider these differences in our estimation.

4.2 Diagnostics

How well do the predictions of the baseline model under Assumptions 1 and 2 agree with actual spending patterns in the data?

The prediction in Proposition 1(ii)—that spending is independent of popularity—cannot be tested because publicly available polling data are too sparse.[19] So we proceed to investigate the predictions of Proposition 2. These predictions are the equal spending ratio result and the constant spending growth result.

Equal Spending Ratios. In Table 2 we look at the extent to which the equal spending ratio result is violated in our data. Since spending ratios are defined as the shares of leftover (rather than total) budgets spent, these ratios can take any value between 0 and 1 every week prior to the final week, where, by construction, they equal 100%. So to not bias the results in the direction of fewer and smaller violations of the equal spending ratio result, we exclude this final week from our analysis.

Table 2 reports that the candidates' weekly spending ratios are within 10 percentage points (pp) of each others' in 80% of election-weeks, and within 5 pp of each others' in 56% (see Table 1 in the Online Appendix for disaggregations of the 5pp analysis). Even in the final six weeks of the campaign when candidates spend larger amounts, they are within 10 pp of each others' in 75% of election-weeks, and within 5 pp of each others' in about half.

Violations of the equal spending ratio result do not seem to be more pronounced in open-seat elections, nor in those where voters are able to cast their ballots early without an excuse. This last finding is consistent with our early voting extension in which the equal spending ratio result continues to hold analytically. On the other hand, we do see more pronounced violations in elections that are lopsided in terms of money spent and final vote shares. If these elections are those in which one candidate (e.g. the better-resourced one) has large leads against the other, then these more pronounced violations could be explained by the variant of our model in which candidates value money left over.

[19]To the best of our knowledge, FiveThirtyEight and Pollster provide the largest publicly available database on polls. We collected data from these sources and identified only 24 elections (all state-wide races) with more than 3 weeks of polling data, which constitutes a sample that is too sparse and potentially not representative of the full set of races in our dataset to conduct a systematic analysis of how spending decisions are affected by changes in relative popularity.

Table 2: $x_t/X_t - y_t/Y_t$

Week	-11	-10	-9	-8	-7	-6	-5	-4	-3	-2	-1
% ∈ (-0.1,0.1)	0.963	0.953	0.938	0.902	0.879	0.847	0.829	0.754	0.676	0.622	0.797
Senate	0.943	0.934	0.975	0.926	0.934	0.885	0.844	0.787	0.746	0.648	0.803
Governor	0.932	0.910	0.887	0.820	0.812	0.812	0.767	0.774	0.639	0.624	0.782
House	0.983	0.977	0.945	0.925	0.884	0.847	0.847	0.734	0.665	0.613	0.801
Early Voting	0.970	0.955	0.942	0.912	0.884	0.844	0.816	0.753	0.673	0.612	0.798
No Early Voting	0.951	0.951	0.931	0.882	0.868	0.853	0.853	0.755	0.681	0.642	0.794
Open Seat	0.942	0.933	0.920	0.897	0.857	0.862	0.866	0.795	0.705	0.656	0.804
Incumbent Competing	0.976	0.966	0.950	0.905	0.891	0.838	0.806	0.729	0.658	0.602	0.793
Close Election	0.976	0.965	0.935	0.941	0.947	0.924	0.906	0.882	0.776	0.706	0.788
Not Close Election	0.958	0.949	0.940	0.886	0.852	0.817	0.798	0.703	0.636	0.589	0.800
Close Budgets	0.974	0.974	0.959	0.925	0.914	0.895	0.883	0.812	0.763	0.695	0.838
Not Close Budgets	0.955	0.937	0.922	0.884	0.851	0.809	0.785	0.707	0.606	0.564	0.764
% ∈ (-0.05,0.05)	0.865	0.815	0.757	0.727	0.661	0.599	0.554	0.468	0.418	0.369	0.562
Average x_t/X_t	0.021	0.028	0.039	0.054	0.075	0.109	0.134	0.184	0.251	0.377	0.728
	(0.032)	(0.036)	(0.044)	(0.051)	(0.054)	(0.067)	(0.073)	(0.085)	(0.095)	(0.108)	(0.076)
Average y_t/Y_t	0.021	0.029	0.038	0.049	0.074	0.105	0.133	0.184	0.249	0.380	0.733
	(0.035)	(0.041)	(0.046)	(0.053)	(0.063)	(0.073)	(0.080)	(0.094)	(0.097)	(0.111)	(0.073)

Note: The table reports the share of elections in which the candidates' spending ratios are within 10 percentage points (or 5 percentage points) of each other for every week, across election types. See the note below Table 1 for definitions of close elections and close budget elections.

Finally, the extent to which our equal spending ratio result appears violated in the data is increasing as the election approaches. One reason for this could be that as election day approaches, spending decisions are more affected by disturbances resulting from factors outside our baseline model.[20] Another possibility is that spending ratios are more likely to be close in percentage points in the early weeks when both candidates spend lower shares of their available budget. To address these possibilities, in the Online Appendix, we also examine the percentage (as opposed to percentage points) differences between the spending ratios of the candidates across weeks. We find that percentage differences tend to decrease (i.e., the equal spending ratio result tends to hold more often) as election day approaches and candidates spend larger amounts.

[20] One such factor is an "October surprise"—the surfacing of new information, like a scandal that creates a wedge between a candidate's forecasted budget (on which some past spending was based) and the budget that actually becomes available. Another factor outside our model is the idea that close to election day, trailing candidates may simply give up because of threshold effects.

Constant Spending Growth. The consecutive period spending ratio (CPSR) is x_{t+1}/x_t for the Democrat and y_{t+1}/y_t for the Republican candidate. In our eleven-week dataset, these variables are defined for ten consecutive week pairs. If the constant spending growth prediction holds, these two ratios should be relatively stable over time. However, since there are candidates who spend zero in some of the earlier weeks, the CPSR cannot be calculated for certain periods. In what follows, we thus calculate CPSRs using two approaches: (i) dropping all elections with zero spending in any week, and (ii) dropping all pairs of consecutive weeks that would include a week with zero spending. Approach (i) leaves us with only 221 (out of the total 601) elections where no zero spending occurs. In approach (ii), instead, we drop 1,692 consecutive week pairs out of a total of 13,222, which is only 12.8%. We also note that in our data there is no instance of zero spending following positive spending: once a candidate starts spending, she continues to do so until the election.

Figure 3 reports the distribution of average CPSRs for every candidate, along with the intervals centered at these averages and width equal to ± 1.96 times the estimated standard deviation. (In the Online Appendix we also report similar plots with the interval defined by the second lowest and the second highest observation for each election.) The distributions obtained from approaches (i) and (ii) are very similar. The reported CPSRs for approach (ii) can be interpreted as growth rates conditional on having started positive spending during an electoral campaign. Approach (ii) discards less data and so we proceed with analyzing the growth rates obtained using such an approach. Hereafter, when we say "growth rates," we refer to growth rates conditional on having started spending positive amounts.

Our baseline model predicts a positive and constant spending growth rate. Looking at Figure 3, the middle 90% of the distribution of CPSR values (5th to 95th percentile) spans $[0.98, 1.9]$. For the candidate with the median value, the average CPSR is 1.16, meaning that her spending increases by 16% on average every week after she starts spending positive amounts. We also find that spending increases from one week to the next for 85% of candidate-weeks. The median standard deviation in candidate CPSRs within an election is 0.814 and more than 75% of candidates have a standard deviation below 2. Variation in CPSR values within an election is typically driven by only a few weeks of volatile growth, rather than by extreme volatility in the entire spending path. Table OA3 in the Online Appendix provides a measure of how the central tendency of candidates' CPSRs within elections varies week by week.

Overall, CPSRs vary within elections, contrary to what our baseline model predicts. One possible explanation for this variation is given by our early voting model in which

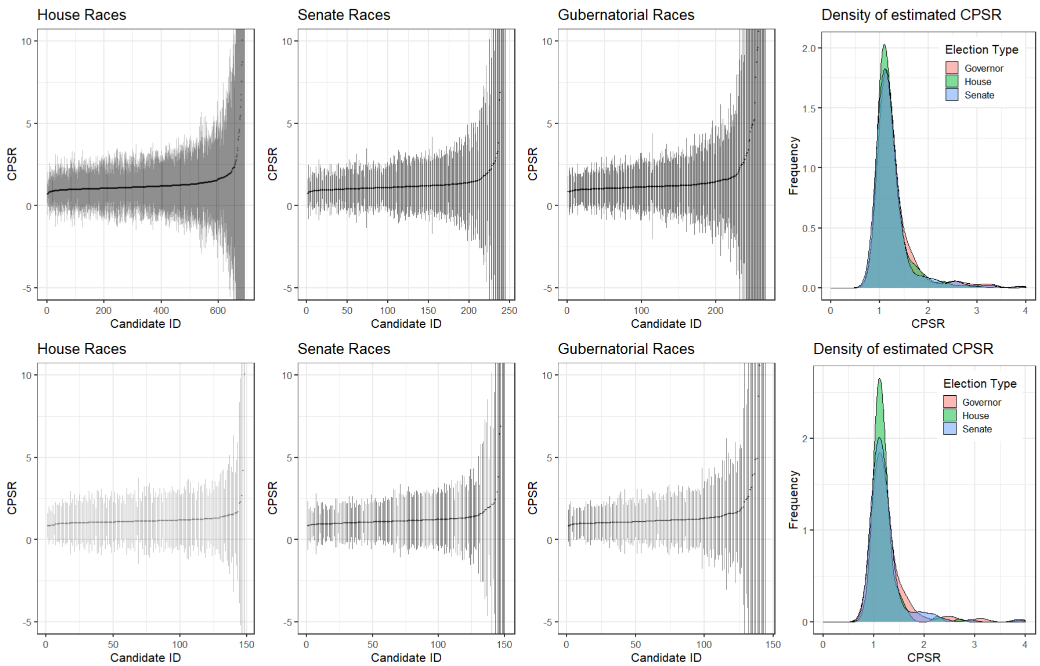

Figure 3: Average CPSR values for candidates in our dataset, along with the interval $[\mu_r - 1.96\sigma_r, \mu_r + 1.96\sigma_r]$, where μ_r and σ_r are the sample average and sample standard deviation of CPSRs. The upper display row depicts the averages that we get by dropping all elections with zero spending. The bottom depicts the averages that we get by dropping all pairs of consecutive weeks that include zero spending. In the first three charts of each row, candidates are sorted based on their average CPSR from lowest to highest. The last chart of each row depicts the densities of average CPSRs across election types from each approach.

spending growth is constant until the time early voting starts, which is typically anywhere from a few days prior to the election to up to eight weeks from election day. Early voting, however, does not appear to be a major driver of violations to the constant spending growth prediction (see Table OA3 in the Online Appendix). Another possible explanation for the deviations from constant spending growth is that candidates value money left over, as in our extension. Though we cannot directly test this, we can reason that if House candidates are more likely to value money left over than Senate or gubernatorial candidates (because the value of office is lower, or their future political ambitions—perhaps to become Senators or governors—are greater, or because they compete more frequently in future elections), this does not appear to be reflected in the disaggregation by election type (again see Table OA3). A third possibility is that the candidates have uncertain budgets that react to their polling performance, as in the evolving budgets model that we present in the Online Appendix. We cannot investigate whether this model can account for the violations from the constant

CPSRs prediction because data on when candidates receive money or pledges from donors are not available.

4.3 Perceived Decay Rates

In our model, the decay rate in popularity leads is $1-\delta$. The perceived decay rate is the value of $1-\delta$ that is "most consistent" with the candidates' spending behavior in an election given that the CPSR in the baseline model is $r=\delta^{1/(\beta-1)}$. Since the perceived decay rate cannot be separately identified from the parameter β using spending data alone, we fix a grid of values of β ranging from 0 to 1 and we report how the distribution of estimated perceived decay rates varies with β.

A straightforward way to estimate the perceived decay rate $1-\delta_j$ in election j, is to let r_j be the mean of the candidates' CPSRs estimated from their actual spending levels in election j (these are given in Figure 3) and then use the relationship $1-\delta_j=1-(r_j)^{\beta-1}$. We perform this estimation using approach (ii) above, namely dropping all candidate-weeks with zero spending. More specifically, δ_j can be estimated directly from the first moment of the distribution of observed CPSRs. Denote

$$r_{j,i,t}=\frac{i\text{'s spending in week } t+1}{i\text{'s spending in week } t},\quad \text{in election } j$$

which is observed for $t=0,1,...,T-2$, for both candidates $i=1,2$ running in election j and can be calculated so long as the candidate spends a positive amount in week t. We compute the first moment of these CPSRs for election j as

$$\hat{r}_j=\frac{1}{|\mathcal{T}|}\sum_{i=1,2}\sum_{t\in\mathcal{T}}r_{j,i,t}$$

where $\mathcal{T}$ is the set of candidate-weeks in election j for which $r_{j,i,t}$ can be computed.[21] Then, as our model predicts $r_{j,i,t}=r_j=(\delta_j)^{1/(\beta-1)}$ for both i and all t, we fix β to some value and estimate the perceived decay rate $1-\delta_j$ from $\hat{r}_j$ as $1-(\hat{r}_j)^{\beta-1}$.

The reason we pool the two candidates' CPSRs to estimate a common perceived decay rate is that this approach increases the precision of our estimates, as it gives us potentially up to 20 total CPSR values (which occurs when there are no weeks with zero spending). In the Online Appendix, we also report candidate-specific decay rates obtained without pooling together the CPSRs of the two candidates. The densities of the estimates we

[21] For example, if both candidates spend positive amounts in all eleven weeks prior to election day, then we have $|\mathcal{T}|=20$.

References

Albert, Z. (2017): "Trends in Campaign Financing, 1980-2016," *Report for the Campaign Finance Task Force, Bipartisan Policy Center. Retrieved from https://bipartisanpolicy.org/wp-content/uploads/2018/01/Trends-in-Campaign-Financing-1980-2016.-Zachary-Albert.pdf.*

Ali, S. and N. Kartik (2012): "Herding with collective preferences," *Economic Theory*, 51, 601–626.

Aravindakshan, A. and P. A. Naik (2015): "Understanding the memory effects in pulsing advertising," *Operations Research*, 63, 35–47.

Bouton, L., J. Cagé, E. Dewitte, and V. Pons (2022a): "Small Campaign Donors," *Working paper.*

Bouton, L., M. Castanheira, and A. Drazen (2022b): "A Theory of Small Campaign Contributions," *NBER Working Paper No. 24413.*

Bronnenberg, B. J., J.-P. H. Dubé, and M. Gentzkow (2012): "The evolution of brand preferences: Evidence from consumer migration," *American Economic Review*, 102, 2472–2508.

Callander, S. (2007): "Bandwagons and Momentum in Sequential Voting," *The Review of Economic Studies*, 74, 653–684.

Chapter 5 of Title 47 of the United States Code 315, Subchapter III, Part 1, Section 315 (1934): "Candidates for public office," `https://www.law.cornell.edu/uscode/text/47/315`.

Choi, J. P., S. M. Chowdhury, and J. Kim (2016): "Group contests with internal conflict and power asymmetry," *The Scandinavian Journal of Economics*, 118, 816–840.

Chung, D. J. and L. Zhang (2015): "Selling to a Moving Target: Dynamic Marketing Effects in US Presidential Elections," *Harvard Business School working paper series# 15-095.*

Coate, S. (2004): "Political Competition with Campaign Contributions and Informative Advertising," *Journal of the European Economic Association*, 2, 772–804.