Expressive Politics: How Animus and Cognitive Dissonance Affect Electoral Extremism

William Howell* Stefan Krasa† Mattias Polborn‡

May 21, 2025

Abstract

We study a model of electoral competition that incorporates two expressive benefits of candidate position taking: the psychological costs of deviating from one's own preferred policy and the psychological benefits of antagonizing an out-group. Whereas concerns about cognitive dissonance consistently temper candidate extremism, the effects of animus are non-monotonic—exacerbating policy divisions when baseline levels are low, and triggering one candidate's capitulation (as distinct from both candidates' moderation) when they are high. We further show that when communication channels are siloed and voters are less concerned about voting for someone who represents their own policy views, candidates are especially inclined to stoke inter-group animosities. Our findings have broad implications for understandings of political polarization, partisan sorting and representation, fragmented media markets, and separation of powers.

*Department of Political Science, University of Chicago. Mailing address: 5828 S. University Ave, Chicago, IL 60637. E-mail: whowell@uchicago.edu.

†Department of Economics, University of Illinois. Mailing address: 1407 W. Gregory Dr., Urbana, IL, 61801. E-mail: skrasa@uiuc.edu.

‡Department of Economics and Department of Political Science, Vanderbilt University, and Department of Economics, University of Cologne. Mailing address: 2301 Vanderbilt Place, Nashville, TN 37235. E-mail: mattias.polborn@vanderbilt.edu.

"Now people vote for candidates not based on what they'll do for them but their ability to inflame the other side which they hate so much." Scott Galloway on Pivot, December 6, 2022.

Electoral politics feature more than just narrow disagreements about policy. They also are infused with social and political identities, psychological needs and wants, and demonstrations of inter-group animosities—or what Brennan (2011) calls "expressive considerations." Elections, thus understood, do more than just alter the direction of public policy-making. They also give voice to voters' self-understandings and feelings about others.

What implications do these expressive considerations have for candidate position-taking and government lawmaking? And quite apart from the policies they choose to endorse, when will candidates deliberately inflame inter-group enmities and thereby amplify the political relevance of these expressive considerations? To investigate these matters, we study a model of electoral competition that incorporates two types of expressive benefits: one that reflects the reputational and/or psychological costs of deviating from one's own preferred policy; and another that captures the benefits of antagonizing an out-group that one does not merely disagree with, but that one actively dislikes. Just as voters wish to minimize their own cognitive dissonance, they occasionally derive pleasure in provoking their opponents—something that, for some supporters of Donald Trump, takes the form of "rage farming"[1] or "owning the libs and pissing off the media... That's what we believe in now. There's really not much more to it."[2]

Expressive considerations, we show, have divergent and often surprising effects on the positions that candidates assume. Whereas previous research shows that a unyielding commitment to policy convictions is associated with political extremism within the electorate (see, e.g., (Zmigrod, 2020)), we find that it creates incentives for political candidates to moderate their own policy positions. Voters who are particularly averse to cognitive dissonance (and hence are reluctant to compromise

[1] As quoted in Erin Douglas, "Texas GOP's voting meme show how Trump-style messaging wins internet's attention." *The Texas Tribune*, January 8, 2012.

[2] As quoted in Tim Alberta, "The Grand Old Meltdown: What Happens When a Party Gives Up on Ideas?" *Politico*, August 24, 2020.

their own policy beliefs) do not tether candidates to extreme and opposing positions. Quite the opposite, these voters induce candidates from competing parties to mitigate their differences.

The effects of animus on equilibrium candidate positions are more subtle but just as remarkable. Increasing animus in both groups leads candidates to diverge. And when these higher levels of animus are concentrated within only one group, the terms of electoral competition abruptly change: while the candidate associated with the subset of voters experiencing high levels of animus stands firm, the opposing candidate abandons the voters with whom she is aligned in order to placate the more aggrieved population. In terms of winning probability, this is successful as she is now, in equilibrium, more likely to win than her opponent, but in terms of policy, this capitulation leaves the temperate portion of the electorate devoid of meaningful political representation.

The model generates new insights in the relevance of non-policy considerations for voting behavior. In standard models in which voters receive a non-policy payoff if a member of their own party wins an election (see, e.g., Adams and Merrill 2003; Erikson and Romero 1990), partisanship matters most when candidates assume the same policy positions; instead, when policy differences between candidates are large, partisanship plays a smaller role because, for most people, the policy utility differences overwhelm any payoff associated with voting for a co-partisan. In our model that operationalizes animus in lieu of party attachment, these dynamics shift in important respects. If differences between candidate positions is large, then one of the candidates is significantly worse for the out-group, and many in-group voters are willing to support that candidate even at a cost of some policy utility loss. As a result, non-policy considerations are activated when candidates diverge.

The model also reveals how obstacles to public policymaking exacerbate polarization. When a candidate's policy proposals are unlikely to become law—either because of legislative gridlock, political dysfunction, partisan intransigence, or institutional weaknesses of the office she seeks—voters may have fewer qualms with voting for someone who does not represent their policy interests. Candidates then have fewer incentives to moderate and deliver, instead, proposals that directly satisfy members of one group by antagonizing the members of the other. The model, as

such, explains an essential aspect of Trump's enduring political appeal—as many voters remain convinced that the political system is irredeemably broken but that Trump, at least, gives voice to their anger towards others, be they Democrats, immigrants, racial minorities, or members of the D.C. establishment (Howell and Moe, 2020; Sides, Tesler and Vavreck, 2018; Webster, 2020).

Finally, the model clarifies the incentives of candidates to foment inter-group hostilities. When channels of communication remain broad and inclusive, we find in an extension that endogenizes animus, political candidates counsel mutual tolerance. Doing so, after all, mollifies voters' appetite for policy extremism. But this rather salubrious result quickly falls apart when communication becomes siloed and candidates can rile up their own group without directly affecting their opponent's. While both candidates would be better off if the animus of both groups was kept in check, each candidate individually benefits from stoking anger within her own affiliated group–particularly when voters assign less importance to the expressive benefits of voting for candidates with whom they agree and when the pathway to delivering on campaign promises remains clear. Lacking any disciplinary mechanism, comity quickly gives way to the lure of demagoguery.

The model, as such, reveals how two foundational expressive needs affect the levels of political polarization that have come to define contemporary American politics. Cognitive dissonance, we find, serves as a moderating force and keeps candidates aligned with their own preferences and those of their core voters. Animus, though, wreaks various forms of havoc. In some instances, it encourages political polarization well beyond what candidates themselves would otherwise prefer. In other instances, it encourages candidates to abandon those voters who harbor no animus whatsoever, and to cater to those groups that are more antagonistic towards their fellow citizens. And though policy-minded candidates would prefer a more accepting and understanding electorate, they themselves confront electoral incentives to stoke anger, antagonism, and division.

We proceed as follows. After the first section summarizes the formal and empirical literatures on expressive voting, the second section introduces the model. The third section distinguishes the group sorting that arises in our model from that observed in standard models of partisan attachment. The fourth section characterizes the equilibria and comparative statics for different values of

animus, and the fifth endogenizes its selection. The sixth section reflects upon the model's implications for partisan polarization, partisan sorting and representation, political communication, and separation of powers. The final section concludes.

1 Literature Review

Our model draws from and contributes to two bodies of research: a formal literature on expressive voting and a predominantly empirical literature on the psychological foundations of political behavior.

1.1 Expressive Voting

Over the last quarter century, a substantial amount of research has been devoted to "expressive voting," or what Clark and Lee (2016) call "the emotional and/or moral satisfaction" that comes from participating in elections (for reviews, see Hamlin and Jennings 2011; 2018). By voting, this literature postulates, people do more than just improve the odds that their favored policy is implemented. They also give voice to their feelings about themselves and others, which satiates a variety of psychological appetites for self-expression (Brennan, 2008).

By attempting to explain why rational actors would pay any cost to vote in the face of vanishingly small probabilities that their ballots will alter the outcome of an election, much of this literature focuses on turnout (Riker and Ordeshook, 1968; Fiorina, 1976; Jones and Hudson, 2000). The mere act of participation, these scholars suggest, yields expressive benefits that may compensate for the inconveniences of voting. Participation, however, is hardly guaranteed. When these expressive needs are not acutely felt, voters may opt to stay at home (Brennan, 2008; Schuessler, 2000). And when candidates' positions diverge sufficiently from those of voters, even when one candidate is strictly preferred over another, alienation may set in and voters may be inclined to abstain (for examples, see Hinich, Ledyard and Ordeshook (1972); Adams, Dow and Merrill (2006)).

Expressive considerations, however, do more than just convince people to show up on Election Day. They also inform the votes that people actually cast. For instance, in Callander and Wilson's (2006; 2008) model of context-dependent voting, a voter's choice between two alternatives depends on the availability of other candidates in the choice set, such that, for example, the presence of a third candidate who opposes immigration may raise the salience of immigration policy differences between the two main candidates. Other models seek to explain how instrumental and expressive considerations jointly translate into vote choices, whether by reference to people's separate preferences for each (Brennan and Hamlin, 1998; Kamenica and Brad, 2014; Taylor, 2015) or the levels of popular support that different candidates from different groups receive within an electorate (Schuessler, 2000).[3]

Three features of this literature warrant emphasis, the first of which concerns the sheer capaciousness of expressive considerations. Scholars recognize all manner of psychological phenomena as the potential subject of expressive voting, ranging from the joys of cheerleading to the pleasures of lambasting an out-group. As Hamlin and Jennings (2011, 333) put it, expressive considerations could include any "aspect of the voter's beliefs, values, ideology, identity or personality regardless of any impact that the vote has on the outcome of the election." In many papers, therefore, expressive benefits are treated as a separate but undifferentiated category of voter preferences (Fiorina, 1976; Jones and Hudson, 2000). And even when specific interpretations are offered, they routinely collapse to a single parameter in a model of turnout or voting (Brennan, 2008; Kamenica and Brad, 2014).

Second, this literature does not characterize how different expressive preferences jointly inform the actions that candidates take. To vote at all, or to vote for a particular candidate, reliably yields such benefits. Expressive benefits, as such, are not endogenously generated through the interactions of voters and candidates. Rather, they are presumed to come with the political territory.

Lastly, there is the referent of these expressive benefits. For most studies, it is the voter herself, as benefits flow from the fulfillment of her own patriotism, sense of civic duty, ideological con-

[3]For a related literature that focuses on labor strikes, see Glazer (1992).

sistency, or some other cognitive need or disposition. A handful of studies, meanwhile, recognize the benefits of expressing views about others, such as the enmity one feels towards an out-group (Glazer, 2008). To our knowledge, however, no one in this literature examines the tradeoffs between different classes of expressive benefits.

1.2 Cognitive Dissonance and Animus

Though certainly not exhaustive, two expressive needs routinely inform our politics: one that affirms a person's own ideological consistency, and another that affords an opportunity to distinguish oneself from the opposition. A substantial amount of empirical work has been conducted on the psychological origins and manifestations of both of these expressive needs.

Going back decades, psychologists have recognized the benefits of signaling one's ideological purity; or, observationally equivalently, reducing the psychological burdens of inconsistency.[4] When making choices of any kind, very much including political ones, people seek to minimize their cognitive dissonance. Simultaneously holding in mind two contradictory thoughts is cognitively taxing, and so too is acting in ways that expressively violate one's core convictions. In electoral politics, consequently, the costs of voting for someone with whom one disagrees appear twice over: first, when pivotal, in the policy losses that may accompany her election; and second, regardless of whether one's vote actually makes a difference in the electoral outcome, in the cognitive dissonance that comes from acting in ways that violate one's policy preferences or principles (Beasley and Joslyn (2001)).

Expressive considerations, however, are not exclusively about self-care. They also encourage attacks on a perceived out-group that one does not merely disagree with, but that one actively dislikes (Webster, 2020). Recent studies on "affective polarization" lay out the basic argument as it relates to Democrats and Republicans in American politics (for a review, see Iyengar et al. (2019)).[5] Rooted in social identity theory, this literature builds upon Henri Tajfel's famous ob-

[4]For a review of much of the field's early development, see Harmon-Jones and Mills (1999).

[5]For related studies in psychology, see (Pietraszewski et al., 2015; Pietraszewski, Cosmides and Tooby, 2014).

servation that members of an in-group will discriminate against an out-group "even if there is no reason for it in terms of the individual's own interest" (Tajfel (1970), 99). Tajfel recognized that discrimination is "extraordinarily easy to trigger" even when groups are randomly assigned (102). In politics, however, groups and allegiances are hardly random. Rather, studies of affective polarization emphasize, political parties offer powerful and salient social identities for many Americans (Mason, 2015; Abramowitz and Webster, 2016); so much so, in fact, that these "mega identities" have become wellsprings of partisan animus in contemporary American politics (Mason, 2018*a*,*b*; Iyengar, Sood and Lelkes, 2012; Iyengar and Westwood, 2015).

It isn't difficult to see how inter-group animosities can whet a person's appetite for antagonism—a character trait that, along with agreeableness, defines one of the five main dimensions of human personality (Lynam and Miller, 2019). Members of an out-group, after all, are often perceived as not merely mistaken or wrong, but as inferior, immoral, or wicked. Compromising with them, as such, invites scorn (Davis, 2019), whereas antagonizing them is cause for minor celebration. Recall, then, the insignia "I really don't care, do u?" written across a green jacket that Melania Trump famously wore in 2019 when she toured an immigration detention center holding children who had been separated from their parents. The first lady shrugged off the public firestorm around the sartorial selection. "I'm driving liberals crazy... You know what? They deserve it."[6] Attuned to her own expressive needs, the First Lady relished the opportunity to antagonize her husband's political adversaries. Their outrage was her delight. As Adam Serwer put it in an *Atlantic* essay, "the cruelty is the point."[7]

And for work on the centrality of anger as a mobilizing force in American politics, see (Valentino et al., 2011).

[6] As quoted in Wolkoff 2020.

[7] Adam Serwer, "The Cruelty is the Point: President Trump and his supporters find community by rejoicing in the suffering of those they hate and fear." *Atlantic*, October 3, 2018. Nor do these dynamics appear to be confined to the Trump presidency. In another more recent *Atlantic* essay on the Republican Party's communication strategies during the Biden Administration, Elizabeth Bruenig observed that "liberal hysteria is no longer an obstacle to good policy making or even an irritating by-product of the democratic process, but rather the desired outcome of almost all right-wing political rhetoric." ("Lauren Boebert's Gun Photo Is Doing Exactly What She Wanted." *Atlantic*, December 8, 2021.) Or as Molly Jong-Fast observes, a Republican Party that remains in the thrall of Trump continues to make "a point of eschewing policy in favor of 'owning the libs' to garner likes, retweets, and small-dollar donations." ("Owning the Libs Is the Only GOP Platform." *Atlantic*, January 12, 2022.)

2 The Model

We envision an electorate with two distinct groups, $i = 1, 2$, that can be understood by reference to either their partisanship (such as Democrat or Republican) or any other salient ascriptive characteristic (such as their race, religion, or language).[8] A citizen in each of these two groups is characterized by an ideal point $\theta \in \mathbb{R}$.

The distribution of ideal points is determined by a stochastic shift-variable M, drawn from a distribution $F(\cdot)$ that is symmetric around zero, and has a non-decreasing (non-increasing) density to the left (right) of zero. A positive (negative) realization of M denotes a shift of all ideal points to the right (left). Formally, for any realization m, the distribution of θ in group i is given by $\Phi_i(\theta - m - \mu_i)$, where Φ_i is symmetric around zero. From an ex-ante perspective (i.e., before M is realized), μ_i is the expected median policy ideal point of group i.

Let ϕ_i and f be the pdfs of Φ_i and F, respectively. To guarantee that second-order conditions for the candidates' optimization problems are satisfied, we assume that the distribution F has a strictly increasing hazard rate on its support, i.e., $f(m)/(1 - F(m))$ is strictly increasing in m (when $F(m) \in (0, 1)$).[9] Finally, let q_i denote the fraction of the population that is of type i.

All voters in our model participate in the election (i.e., no abstention), and their only choice problem is who to vote for. Consistent with Hamlin and Jenning's (2011, p. 650) observation that "expressive and instrumental motivations are best seen as joint inputs into an overall analysis of behavior," our model incorporates a richer set of voter considerations than the canonical spatial voting model. Specifically, the utility of a voter θ in group i from voting for a candidate with proposed policy x is

$$u_{\theta,i}(x) = \alpha_i|x - \mu_{-i}| - \beta|x - \theta|, \tag{1}$$

[8] In principle, we could consider a more general setting with multiple groups that experience varying levels of animus toward one another (including, possibly, some groups that are neutral and others that either are just targets or just sources of animus). It is not obvious, however, that an analysis of this considerably more complex environment would yield qualitatively new insights beyond what is recovered in the simpler setting that we study with just two groups.

[9] This is a standard assumption that is satisfied, for example, for uniform and normal distributions.

where the first term is the main innovation. Dropping it, we return to the standard Calvert-Wittman model of electoral competition. But consistent with the literature on affective polarization, its inclusion recognizes that voters are motivated, to some degree, by a political dislike of the out-group, denoted $-i$. We further assume that this animus can be satisfied by voting for a candidate who espouses a position detrimental to their interests.

Specifically, we assume that the policy disutility of the "typical" out-group member satisfies a voter's appetite for antagonisms and thereby increases her utility. One might imagine the intended target is someone other than the median member of the out-group, including, perhaps, all members of the out-group (which would require integrating over all of their types) or particularly extreme members of the out-group (for example, when thinking about the effects of student loan forgiveness, Republicans might think more of a purple-haired gender-studies major rather than a more typical Democratic-leaning student). Our selection of the median out-group member allows for relatively simpler analyses that produce substantively comparable results.

The weight on the animus utility is denoted α_i and setting $\alpha = 0$ reduces our model to the typical case in the literature where voters care only about policy in relation to their own ideal point. The core findings in Section 4 do not depend on using linear-loss functions in (1). However, with a quadratic loss function, we can only establish the existence of local equilibria, which would generate technical issues when we endogenize animus in Section 5.

The second term corresponds to a standard spatial policy utility, though we assume that this part of the voter's payoff is also expressive (that is, people *like to vote* for the candidate whose policy they prefer) rather than based on pivotality considerations (i.e., they only care who they vote for if it influences who wins).[10] Interpreted this way, the first term focuses on animus, the second term measures cognitive dissonance, and α_i/β represents the marginal rate of substitution

[10]Observe that models where voters are only motivated by instrumental considerations (e.g., costly voting models) cannot explain the large participation rates that we routinely see in large elections. In contrast, our model of expressive voting is consistent with the positive participation rates, even when the election outcome is not in question. A further indication of the importance of the expressive utility component is that, in elections with multiple candidates, a significant number of voters cast their ballots in favor of candidates who have no chance of winning ("sincere voting").

at which voters in group i trade off their two expressive needs.[11]

We assume that $\alpha_i/\beta < 1$ for both groups. Substantively, this implies that all voters have single-peaked preferences.[12] How should we think about the determinants of the relative magnitudes of α_i and β? Clearly, α_i reflects to the degree to which group i dislikes the out-group. The parameter β, meanwhile, depends on the cost of "cognitive dissonance" for an individual who votes for a candidate who espouses a policy position that differs from her own.

One particularly important consideration regarding β is as follows. Suppose that the winning candidate only succeeds in implementing a policy platform with probability p, while with probability $1 - p$, the status quo prevails. The costs of cognitive dissonance may then depend on the expectation of having to live under a policy that one voted for. Thus, when a political system is mired in gridlock and the policy platforms of the winning candidate are unlikely to be implemented (i.e., p is small), this cost is small and β is low. By contrast, when policy proposals reliably become law, these costs are likely to rise and β increases in value.

The two candidates' ideal points θ_L and θ_R are (without loss of generality) normalized to be symmetric around zero, $\theta_L = -\theta_R$. The candidates are entirely policy-motivated, with utility $u_P(x) = -|x - \theta_P|$, $P \in \{L, R\}$, where x is the implemented platform.[13] While candidates themselves harbor no animus against any group, they are aware that voters do, and thus rationally consider voters' reactions when they (simultaneously) choose their policy positions $x_L, x_R \in \mathbb{R}$.

[11] Other models attend to the costs of cognitive dissonance, thus understood, in different ways. Rather than trading off competing expressive needs within a political environment in which ideal points are fixed, voters in (Callander and Carbajal, 2022) dynamically update their ideal points toward the candidates for whom they vote.

[12] If $\beta < \alpha$, a voter is happiest with policies that are very far away from his "ideal" point because they make the out-group suffer as well. Such an extreme level of animus is probably not realistic. The interested reader can find an analysis of this case in a working paper version of this paper.

[13] Of course, candidates themselves may take certain pleasure in antagonizing an out-group or from simply holding public office. By characterizing these candidates as strictly policy motivated, we examine a hard case for the proposition that expressive voting considerations induce changes in candidate position-taking. If candidates themselves feel animus against the out-group, our main results are only strengthened. Meanwhile, the inclusion of a payoff from holding office does not alter any of the main comparative statics that follow.

3 Voter behavior

In most models of electoral competition, candidates face a standard tradeoff: by moderating their policy position, they improve their chances of winning; but conditional on winning, candidates increase their policy utility by adopting a more extreme position. The introduction of different forms of expressive benefits, we find, alters this tradeoff in interesting and non-obvious ways.

To determine how candidates' positions affect their probability of winning, we need to find the cutoff value $m(x_L, x_R)$ for the shift parameter for which the election ends in a tie, given platforms x_L and x_R. The first step towards this objective is to identify and analyze the behavior of indifferent voters. The voter type θ in group i who is indifferent between the candidates is given by

$$\alpha_i|x_L - \mu_j| - \beta|x_L - \theta| = \alpha_i|x_R - \mu_j| - \beta|x_R - \theta|. \tag{2}$$

Note first that, if $x_L = x_R$, then (2) holds for all θ—i.e., all voters are indifferent between the candidates, exactly as in the case with standard voter preferences. Intuitively, the introduction of expressive preferences does not change this result, because they are *not* based on a direct "partisan" preference for candidates "associated" with the voter's in-group or out-group, as in Adams and Merrill (2003) and Erikson and Romero (1990). Rather, expressive preferences work through the candidates' proposed policies, and when both candidates choose the same position, each voter is rendered indifferent.

If, instead, $x_L \neq x_R$, Lemma 1 below (proved in the Appendix, along with all other formal results) analyzes how expressive considerations inform candidate position-taking.

Lemma 1 *Let $x_L < x_R$, and let $\bar{x} = \frac{x_L + x_R}{2}$ be the midpoint between the candidates' positions. Then*

the group i voter who is indifferent between candidates located at x_L and x_R is given by

$$\bar{\theta}_i = \begin{cases} \bar{x} \cdot \frac{\beta - \alpha_i}{\beta} + \frac{\alpha_i}{\beta}\mu_j & \text{if } x_L < \mu_j < x_R; \\ \bar{x} + \frac{\alpha_i}{\beta} \cdot \frac{x_R - x_L}{2} & \text{if } \mu_j \geq x_L, x_R; \\ \bar{x} - \frac{\alpha_i}{\beta} \cdot \frac{x_R - x_L}{2} & \text{if } \mu_j \leq x_L, x_R. \end{cases} \tag{3}$$

To understand the implication of Lemma 1, suppose that $\bar{x} = 0$ and consider Figure 1. The cutoff $\bar{\theta}_1$ for group 1 is to the right of $\bar{x}$, while the reverse is true for cutoff $\bar{\theta}_2$ for group 2. Intuitively, from the perspective of a group 1 voter, Candidate L's policy hurts the other group more than Candidate R's, so that $\bar{\theta}_1 > \bar{x}$.[14] An analogous argument establishes that $\bar{\theta}_2 < 0$. Thus, in contrast to the standard model, ours yields some overlap in the policy preferences of left and right candidate supporters. Note that decreasing α_i/β decreases the gap between $\bar{\theta}_1$ and $\bar{x}$. If animus is zero, then we are in the standard case where the cutoff between candidate L and candidate R supporters is $\bar{x}$.

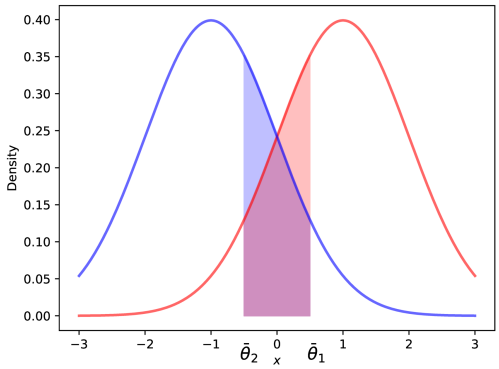

Figure 1: Overlapping voters: Group 1 voters to the left of $\bar{\theta}_1$ vote for candidate L, and group 2 voters to the right of $\bar{\theta}_2$ vote for candidate R.

It is instructive to compare our model with one in which each voter has a *partisan* attachment to one of the parties; that is, on top of the policy payoff, partisans receive an additional non-policy

[14] Algebraically this inequality follows immediately from (3). If x_L and x_R are more extreme, then the first case applies, as $\mu_2 > 0$ the cutoff $\bar{\theta}_1$ is strictly positive. The second case applies when x_L and x_R are more centrist. Again, $\bar{\theta}_1$ is strictly positive because $x_R - x_L > 0$. Further, despite some non-concavities in voter utilities, it is still the case that all voters of group i prefer Candidate R if they are to the right of $\bar{\theta}_i$, and Candidate L if they are to the left $\bar{\theta}_i$.

payoff if a member of their own party wins. In this setting, overlap also arises, because some Democratic partisans who are policy-wise closer to the Republican position will nonetheless vote for a member of their own party, just as Republican partisans do the same.

The forces supporting this overlap, however, are quite different. In the partisan model, partisanship matters most when candidates assume the same policy positions; that is, the size of the overlap in the partisan model reaches its maximum when the voter is policy-wise indifferent between the candidates. In such a setting, almost all Democratic and Republican partisans (irrespective of their policy preferences) vote for their own party's candidate. In terms of the voters' ideological positions, the overlap is complete. But when the policy difference between the parties is large, partisanship plays a smaller role because, for most people, the policy utility difference between the parties outweighs the partisan payoff.

In our model, by contrast, the overlap results from voters' animus toward one another. If the parties' policy positions are indistinguishable from one another, neither candidate provides much of an advantage in antagonizing the out-group. In this case, the cutoffs of both groups are very close to the average party position, and the overlap is minimal. If, instead, the difference between the parties' positions is large, then one of the candidates is significantly worse for the out-group, and many in-group voters are willing to support that candidate even at a cost of some policy utility loss. Of course, the behavior among voters of the other group is symmetric. Thus, in our model, large policy differences between parties result in a large overlap in the two groups' voting behavior.

Knowing the indifferent voter in each group allows us to determine m, the critical value of the shift parameter M that determines which candidate wins the election. The election ends in a tie if

$$q_1\Phi_1\left(\bar{\theta}_1 - m - \mu_1\right) + q_2\Phi_2\left(\bar{\theta}_2 - m - \mu_2\right) = \frac{1}{2}, \tag{4}$$

and the left (right) party wins for smaller (larger) values of M. In Lemma 2, we explicitly solve equation (4) for m when the groups are equal-sized and have preference distributions that have the same shape.

Lemma 2 *If $q_1 = q_2$ and $\Phi_1 = \Phi_2$, then the state at which the election ends in a tie is*

$$m(x_L, x_R) = \frac{\bar{\theta}_1 + \bar{\theta}_2 - \mu_1 - \mu_2}{2}, \tag{5}$$

where $\bar{\theta}_i$ is defined in (3).

Lemmas 1 and 2 together reveal that voter animus reduces candidates' incentives to moderate in electoral competition. To see this crucial result, suppose, e.g., that $\theta_L < x_L < \mu_1 < \mu_2 < x_R < \theta_R$ and hence

$$m(x_L, x_R) = \frac{\alpha_1\mu_2 + \alpha_2\mu_1}{2\beta} + \left[1 - \frac{\alpha_1 + \alpha_2}{2\beta}\right]\frac{x_L + x_R}{2}. \tag{6}$$

Clearly, the higher $\alpha_1 + \alpha_2$, the smaller is a candidate's electoral benefit of moderation (in terms of by how much moderation shifts the indifferent voter). By contrast, the cost of moderation only depends on the candidates' policy preferences and is thus the same as in the standard model.

4 Equilibrium Analysis

In this section, we assume that the groups are of equal size, located symmetrically around zero (i.e., $\mu_2 = -\mu_1 = \mu$), and that the candidates' ideal positions are also symmetric around zero and more extreme than the medians of the two population groups ($-\theta_L = \theta_R > \mu$). Furthermore, the amount of uncertainty about the position shift parameter is intermediate; that is, $\mu < 1/(2f(0)) < \theta_R$.

Candidates trade off the benefits of positioning closer to their respective ideal points against the costs of a reduced winning probability. Formally, Candidates L and R solve, respectively,

$$\max_{x_L} -F(m(x_L, x_R))|x_L - \theta_L| - (1 - F(m(x_L, x_R)))|x_R - \theta_L|, \tag{7}$$

$$\max_{x_R} -F(m(x_L, x_R))|x_L - \theta_R| - (1 - F(m(x_L, x_R)))|x_R - \theta_R|. \tag{8}$$

Let $\bar{\alpha} = (\alpha_1 + \alpha_2)/2$ denote the average level of animus and let $\mu = \mu_2 = -\mu_1$. In Proposition 1

we show that $\mu < 1/(2f(0)) < \theta_R$ implies $x_L < -\mu$ and $x_R > \mu$. Hence equation (6) applies, i.e.,

$$m = \frac{\mu(\alpha_1 - \alpha_2)}{2\beta} + \frac{\beta - \bar{\alpha}}{\beta}\frac{x_L + x_R}{2}. \tag{9}$$

As Lemma 3 in the Appendix shows, for an interior equilibrium with $x_L > \theta_L$ and $x_R < \theta_R$, the following first-order conditions are necessary and sufficient:

$$-f(m(x_L, x_R))\left(\frac{\beta - \bar{\alpha}}{2\beta}\right)(x_L - x_R) - F(m(x_L, x_R)) = 0, \tag{10}$$

$$-f(m(x_L, x_R))\left(\frac{\beta - \bar{\alpha}}{2\beta}\right)(x_R - x_L) + (1 - F(m(x_L, x_R))) = 0. \tag{11}$$

Adding (10) and (11), the first terms of both equations cancel out, implying that $F(m) = 1/2$ (i.e., both candidates win with probability $1/2$) and $m = 0$. Substituting $m = 0$ in (9) and (10) and solving the resulting equation system yields

$$x_L = -\frac{\beta - f(0)\mu\Delta\alpha}{2f(0)\,(\beta - \bar{\alpha})}; \tag{12}$$

$$x_R = \frac{\beta + f(0)\mu\Delta\alpha}{2f(0)\,(\beta - \bar{\alpha})}, \tag{13}$$

where $\Delta\alpha = \alpha_2 - \alpha_1$ is the difference between group 2's animus against group 1, and the animus in the opposite direction.[15]

We now can define elite polarization as the distance between the policy positions of Candidates R and L, which is given by

$$x_R - x_L = \frac{2\beta}{2f(0)\,(\beta - \bar{\alpha})} = \frac{1}{f(0)\left(1 - \frac{\bar{\alpha}}{\beta}\right)}. \tag{14}$$

Recall that, if animus $\bar{\alpha} = 0$, then we are in the standard Calvert-Wittman model of policy-

[15] These formulas apply as long as $\theta_L \leq x_L$ and $\theta_R \geq x_R$. If (12) implies a value of x_L that is less than θ_L, then the actual policy of candidate L is at L's ideal point θ_L. The same is true for candidate R. Thus, if animus $\bar{\alpha}$ becomes sufficiently large (but $\bar{\alpha} < \beta$), then both candidates will choose policies that are at their respective ideal points. In the remainder of the discussion in this section, we focus on the case of an interior solution.

motivated candidates where polarization is solely determined by the uncertainty over the location of the median voter. Specifically, if $f(0)$ is smaller (i.e., there is more uncertainty about the location of the decisive voter), then the marginal chance that a candidate can achieve victory by moderating goes down, and thus polarization $x_R - x_L$ increases.

While this effect is also present in our model, it additionally shows that changes in animus affect elite polarization even if uncertainty about the median voter's position remains unchanged. As $\bar{\alpha}$ increases, equation (14) implies that polarization goes up. Increasing animus leads to polarization between candidates. Intuitively, even though the candidates' objectives are directly unaffected by animus, voters in both groups become less willing to switch to the candidate associated with the out-group, and therefore both candidates' incentive to moderate decreases.

Interestingly, polarization depends only on the overall average value of animus $\bar{\alpha}$, but not on the animus difference $\Delta\alpha$. This is because a larger $\Delta\alpha$ (while keeping the average animus $\bar{\alpha}$ constant) moves both x_L and x_R by the same amount to the right, i.e., towards the median ideal point of the group whose relative animus increased.

Equations (12) and (13) imply that increasing the animus of the group on the right against the group on the left results in more extreme platforms. The derivative of x_L with respect to α_2 is

$$-\frac{2\alpha_1 f(0)\mu + \beta(1 - 2f(0)\mu)}{(2\beta - \alpha_1 + \alpha_2)^2 f(0)} < 0,$$

because $f(0) < 1/(2\mu)$ and hence $1 - 2f(0)\mu > 0$. Thus, increasing α_2 shifts x_L to the left. The derivative of x_R with respect to α_2 is

$$\frac{\beta + 2(\beta - \alpha_1)f(0)\mu}{(2\beta - \alpha_1 + \alpha_2)^2 f0} > 0.$$

Hence, increasing α_2 moves x_R to the right. In the Appendix we show that these results also hold for the boundary case where $x_R = \theta_R$ or $x_L = \theta_L$.

Finally, consider the effects of the cognitive dissonance weight, β, on policy choices and polar-

ization. Remember that one of the interpretations was that a higher probability that a candidate's proposed policy would be implemented translates into a higher β.[16]Inspection of (14) shows immediately that a higher β reduces polarization. Since β appears in both numerator and denominator of the policy position formulas in (12) and (13), it is less obvious how each candidate's position changes as β changes. The proof of Proposition 1 shows that the candidate who is associated with the more antagonistic group (i.e., if $\Delta\alpha > 0$, Candidate R) moderates when β increases; for the other candidate, we provide a necessary and sufficient condition for moderation.

Intuitively, an increase in β increases the marginal voter's cost of posturing. Thus, voters react more strongly to candidate positions, which generally encourages candidates to moderate their positions. In particular, this effect dominates for Candidate R (considering the case $\Delta\alpha > 0$).

For Candidate L, an additional effect is that an increase in β increases the winning probability of Candidate L, the candidate with the more moderate position. However, a higher winning probability means that the opportunity cost of (further) moderation increases. Whether this effect or the moderation effect described above dominate for Candidate L depends on parameters.[17]

Proposition 1 summarizes our results so far:

Proposition 1 *Under the assumptions of this section (i.e., $\theta_R = -\theta_L$, $\mu_2 = -\mu_1 = \mu$, $\mu < 1/2f(0) < \theta_R$, $\Delta\alpha \geq 0$ and $\bar{\alpha} \leq \beta$) there exists a unique equilibrium with the following properties:*

1. *Increasing animus α_i of any group $i = 1, 2$ strictly increases elite polarization $x_R - x_L$ until both policies are at the candidates' ideal points θ_L and θ_R, respectively.*

2. *If the animus of the group on the right towards the group on the left increases, then each candidate's platform becomes more extreme, unless the platform is already at the candidate's ideal point.*

3. *Elite polarization, $x_R - x_L$, decreases as β increases.*

[16]For example, an increased probability that a proposed policy does not just irritate an out-group, but that it is actually implemented, would result in a higher weight on the cognitive dissonance cost because it would be costly for the voter to know that he voted for a policy that he dislikes and now has to live with.

[17]For Candidate R, an increase in β decreases his winning probability, so that both effects go in the same direction.

Furthermore, an increase in β definitely moderates Candidate R ($x_R \downarrow$), and moderates Candidate L's position ($x_L \uparrow$) if and only if $\bar{\alpha} > f(0)\Delta\alpha$.

Next, we consider the case where the targets of animus are more extreme than the ideal positions of the respective candidates (formally, we assume that $-\mu < \theta_L$ and $\theta_R < \mu$). In a world where partisan voters often have very unrealistic views about the typical member of the opposing party, this is a relevant case. For $0 \leq \Delta\alpha < 2\beta$, equilibrium policies x_L and x_R are between the candidates' ideal points. Thus, $-\mu < x_L, x_R < \mu$, which in turn implies that voters' utilities are concave on $[-\mu, \mu]$. Thus, we will assume for the following results that $0 \leq \Delta\alpha < 2\beta$ instead of $\alpha_i < \beta$.

If $-\mu < x_L < x_R < \mu$ then (3) and (5) imply that the cutoff voter is given by

$$m(x_L, x_R) = \frac{x_L + x_R}{2} - \frac{\Delta\alpha}{2\beta} \cdot \frac{x_R - x_L}{2}. \tag{15}$$

We show in the Appendix that when $\Delta\alpha$ becomes large, Candidate R chooses his ideal policy $x_R = \theta_R$. Thus, Candidate L's first-order condition becomes

$$f(m(x_L, \theta_R))\left(\frac{1}{2} + \frac{\Delta\alpha}{4\beta}\right)(\theta_R - x_L) - F(m(x_L, \theta_R)) = 0 \tag{16}$$

Substituting $x_L(\Delta\alpha)$ for x_L and taking the derivative of this first-order condition (16) with respect to $\Delta\alpha$ yields

$$x_L'(\Delta\alpha) = \frac{4\beta f(m)(\theta_R - x_L)}{(2\beta + \Delta\alpha)(8\beta f(m) - (2\beta + \Delta\alpha)(\theta_R - x_L)f'(m))}. \tag{17}$$

In the Appendix we show that Candidate L's winning probability $F(m) > 0.5$. Because F is symmetric around zero and single peaked this implies that $f'(m) < 0$. Thus, both the numerator and the denominator of (17) are strictly positive, i.e., x_L increases as animus of the right against the left increases.

Intuitively, Candidate L moves closer to Candidate R's policy position in order to reduce the negative impact of animus against the left's position. This also implies that the polarization be-

tween the political positions of the right and the left candidate, $\theta_R - x_L$, decreases. From the perspective of a single conservative who is not himself affected by antagonism but rather interested only in policy outcomes, strongly asymmetric antagonism (i.e., right-wing voters hating on the left more strongly than in the opposite direction) thus has two effects: On the positive side, the left candidate essentially capitulates in terms of the policy he puts forward in order to escape the negative electoral consequences of animus. On the negative side, the left candidate is more likely to win than the right candidate. However, we can also show that the winning probabilities are constant once the right candidate's policy position has reached θ_R so that, marginally, increased antagonism only has the beneficial effect on the left candidate's policy position.

We summarize these results in Proposition 2. The detailed arguments are in the Appendix.

Proposition 2 *Let $\theta_R = -\theta_L$, $\mu_2 = -\mu_1 = \mu$, and $\mu > \theta_R$. Then there exist $\bar{\Delta} < 2\beta$ such that if animus against the left is sufficiently larger than animus against the right, i.e, $\bar{\Delta} < \Delta < 2\beta$, then :*

1. *Candidate L's winning probability is greater than 0.5.*
2. *Candidate R choose a position equal to his ideal point ($x_R = \theta_R$); increasing animus, $\Delta\alpha$, against the left increases x_L, decreases elite polarization $x_R - x_L$ and leaves the candidates' winning probability unaffected.*

5 Endogenous Animus

In addition to responding to people's pre-existing antipathies, politicians also stoke them.[18] To clarify the conditions that support such activity, we now allow candidates to manipulate the amount of animus one group of voters feels about another (or both feel about each other).

[18]For example, Ash, Morelli and Van Weelden (2017) examine how members of Congress allocate time across different issues in their floor speeches, and they find that US senators focus on divisive issues when they are up for election.

When endogenizing α_i, candidates' behavior depends on whether increases in animus felt by one group can occur without affecting the other; or whether, instead, animus increases in both groups. In the latter case, wherein a candidate can only increase α_1 and α_2 by the same amount, inflammatory speech has no benefit. As Proposition 1 indicates, if $\alpha = \alpha_1 = \alpha_2$ and α is increased but remains below $\tilde{\alpha}$, then both candidate positions become more extreme, but the candidates' ex-ante expected utilities do not change.[19] If α rises above $\tilde{\alpha}$, then both candidates are actually worse off. Consequently, if a candidate pays any cost associated with increasing animus, then $\alpha = 0$ is optimal.

Things look very different, however, when candidates communicate through segregated communication channels. When an increase in the animus felt by one group does not automatically induce equivalent increases felt by the other, candidates may have incentives to foment inter-group hatred—at least under some well-specified conditions. To see this, suppose again that both groups have the same size ($q_1 = q_2 = 1/2$), with $\mu_2 = -\mu_1 = \mu$, and that ϕ and f are symmetric. We consider the setting of Section 4 with symmetric candidates ($\theta_R = -\theta_L$) and $\mu < 0.5/f(0) < \theta_R$, but with two important differences.

First, each candidate can choose, at cost $c(\alpha_i) = b\alpha_i$, the level of α_i in the group with which she is ideologically aligned. After this choice becomes common knowledge, the game proceeds as before with the candidates adopting positions, voters selecting a winner, and proposed policies being implemented probabilistically.

Second, we now explicitly separate the forces that were, thus far, both subsumed into β: first, an explicit dislike of voting for a candidate who espouses positions different from one's own, regardless of whether this policy is implemented; and second, a desire for absolution from voting for a candidate who is ultimately responsible for implementing an objectionable policy. For simplicity, we consider an additive form such that $\beta = \hat{\beta} + p$, where p is the probability that the proposed platform is implemented. With probability $1 - p$, the policy remains at some status quo x_S, which

[19] The latter follows from the fact that our candidates have absolute distance preferences and so are, in effect, risk-neutral with respect to policy lotteries, as long as all realizations of the lottery are on one side of their ideal point. If, instead, candidates have strictly concave utility functions, they would dislike the larger risk even in this case.

is the same regardless of which candidate is elected.

The parameters $\hat{\beta}$ and p have different effects on politicians, who remain entirely policy-motivated. While $\hat{\beta}$, which represents voter preferences, has no direct effect on a policy-motivated politician's utility, p determines how likely a politician's choice of α_i is to yield actual changes in policy. In summary, candidate R's first-period optimization problem is

$$\max_{\alpha_2 \geq 0} -pF(m(\alpha_1, \alpha_2))|x_L(\alpha_1, \alpha_2) - \theta_R| - p(1 - F(m(\alpha_1, \alpha_2)))|x_R(\alpha_1, \alpha_2) - \theta_R| - b\alpha_2, \tag{18}$$

and analogously for candidate L. Voters' behavior in the second period is unchanged from above, given the levels of animus α_1 and α_2 chosen by the candidates.[20]

Remember that, for equilibria with $\theta_L < x_L < x_R < \theta_R$, the policy positions are given by (12) and (13), where we now substitute β by $\hat{\beta} + p$. We first analyze whether there exist equilibrium animus levels α_1 and α_2, such that policies are strictly between θ_L and θ_R. For such an equilibrium to exist, the winning probabilities must be $1/2$ even if $\alpha_1 \neq \alpha_2$. In the first stage, therefore, Candidate R solves

$$\max_{\alpha_2} \frac{px_L(\alpha_1, \alpha_2) + px_R(\alpha_1, \alpha_2)}{2} - b\alpha_2 \text{ s.t. } (1)\ x_R(\alpha_1, \alpha_2) \leq \theta_R,\ (2)\ x_L(\alpha_1, \alpha_2) \leq \theta_L. \tag{19}$$

Suppose that both constraints of problem (19) are slack. Taking the derivative of the objective with respect to α_2 yields $0.5(\beta - \alpha_1)\mu/(\beta - \bar{\alpha})^3 > 0$. Thus, the objective of (19) is strictly convex in α_2. Hence, the solution is either at $\alpha_1 = 0$ or at a point where either constraint (i) or (ii) binds. A similar argument holds for Candidate L.

Proposition 3, whose remaining steps are proved in the Appendix, summarizes the candidates' decisions about whether to stoke animus.

Proposition 3 *Suppose that $\theta_R = -\theta_L$, $\mu_2 = -\mu_1 = \mu$, the distribution and size of groups is identical, and $\mu < 1/(2f(0)) < \theta_R$. Then,*

[20] Note that the utility derived from the status quo, x_S, can be dropped from the optimization problem, because it is not affected by the choices of animus and policies.

1. *in any subgame-perfect and symmetric pure strategy equilibrium, animus is either* $\alpha_1 = \alpha_2 = 0$ *or* $\alpha_1 = \alpha_2 \geq \tilde{\alpha} = (\hat{\beta} + p)\left(1 - \frac{1}{2\theta_R f(0)}\right) > 0$*;*

2. *there exists* $\bar{b} > 0$ *such that a subgame-perfect, pure strategy equilibrium with no animus* ($\alpha_1 = \alpha_2 = 0$) *exists if and only if marginal costs are sufficiently large, i.e.,* $b \geq \bar{b}$*. The cost cutoff* $\bar{b}$ *is strictly increasing in* $\hat{\beta}$ *and strictly decreasing in* p*. In particular, if no animus is an equilibrium for some value of* $\hat{\beta}$ *and* p*, then it remains an equilibrium if* $\hat{\beta}$ *is increased or* p *is decreased.*

Proposition 3 reveals that any increase in animus will occur rapidly if the cost of increasing animus or the candidates' policy preferences fluctuate over time. A very slight change in either of these underlying fundamentals may move the candidates from a situation in which widespread comity prevails ($\alpha_1 = \alpha_2 = 0$) to one in which candidates incite significant inter-group hostilities.

A candidate's interest in stoking animus depends upon $\hat{\beta}$ and p. The second point of Proposition 3 shows that when $\hat{\beta}$ increases, the critical cost threshold beyond which a candidate no longer invests in animus decreases. Intuitively, if $\hat{\beta}$ is large, changes in animus affect voter behavior less, and hence are less valuable to candidates seeking to win election. As a result, when voters are particularly concerned about the costs of cognitive dissonance, candidates choose lower levels of animus.

The second point of Proposition 3 also shows that when p decreases, the critical cost threshold beyond which a candidate no longer invests in animus increases, leaving fewer such cases. To see why, remember that candidates incur a reputational cost for increasing the level of animus; and the only reason candidates are willing to pay this cost is to obtain a policy platform that is closer to their own ideal point. If the probability p of being able to implement one's policy proposals is small, then the cost of stoking inter-group hostilities is no longer worthwhile. The production of animus, as such, arises out of concern for the genuine policy stakes of electoral competition.

6 Discussion

By embedding a richer voter utility function within a standard model of electoral competition, we discover new insights into the origins of partisan polarization, the dynamics of partisan sorting and representation, the animosities that roil our new, more fragmented media landscape, and the electoral consequences of separation of powers.

Partisan Polarization. Our model speaks most directly to an empirical phenomenon that has long puzzled scholars of American politics: namely, why the two major parties have grown increasingly polarized at a time when voters, in the main, have remained ideologically moderate (Ansolabehere, Rodden and Snyder, 2006; Fiorina, Abrams and Pope, 2008; Fowler et al., 2023; Hill and Tausanovitch, 2015).[21] To reconcile these two facts, scholars have offered a variety of explanations that implicate the rise of money in politics (e.g., (Baron, 1994; Moon, 2004; Ensley, 2009)), changes in partisan coalitions (Levendusky, 2009), political activists (Layman et al., 2010), partisan primaries (Hill and Tausanovitch, 2015, 2018; Hirano et al., 2010; Krasa and Polborn, 2018), rule changes within Congress (Theriault, 2008; Polborn and Snyder Jr, 2017), increasing partisan competition (Lee, 2009), and growing wealth and income inequality (McCarty, Poole and Rosenthal, 2016).[22] To the mix, we add another that focuses our attention squarely on the voters' expressive needs.

Expressive considerations are not of a piece. Moreover, we show, they pull in different directions. As voters care more about minimizing cognitive dissonance, our model reveals, candidates assume increasingly moderate positions. But as voters' animosities toward an out-group grow, which is precisely what empirical studies of affective polarization have shown to have happened over the last several decades in the United States (Boxwell, Gentzkow and Shapiro, 2024), candidates generally drift to the extremes. A political world in which anchoring beliefs seem less important and inter-group hostilities become inflamed, our model reveals, supports rising levels of

[21] Note, too, the even longer body of work in comparative politics on May's "law of curvilinear disparity," which suggests that rank-and-file party members will be more extreme than either party leadership or voters (May, 1973).

[22] For further discussion of these and other potential causes of polarization, see Barber and McCarty 2018; McCarty 2019.

polarization.

Notice, too, how the voters themselves drive these changes. Whereas much of the existing literature on the causes of polarization points towards external factors that push against the otherwise moderating influence of voters (for a review, see (McCarty, 2019), chapter 5), and whereas a vast behavioral literature suggests that ill-informed and politically naive voters blindly follow members of their own party or reflexively reject members of the opposition (e.g., (Achen and Bartels, 2017)), our model reveals how policy polarization can result from political elites exploiting divisions among voters that are unrelated to policy. To draw a straight line from polarized elites to an electorate, one need not reject the possibility that voters hold relatively moderate policy views; nor must one adopt an especially dim view of voters' agency or knowledge. Rather, one need only recognize the expressive needs that inform people's voting behavior.

Partisan Sorting and Representation. A substantial body of work investigates when voters sort themselves into parties that best represent their policy preferences; and when, instead, voters support a party with which they merely share some allegiance or identity (see, e.g., Levendusky (2009)). Our model clarifies how expressive considerations affect both the propensity of voters to ally themselves with a party that may not represent their policy views well, and the implications their choices have for candidate behavior.

Animus, we show, reifies partisanship and other ascriptive characteristics. As animus towards an out-party increases, voters are inclined to stick with the party with which they are associated, even when the other party does a better job of representing their policy views. Recognizing these calculations by voters, the candidates tack toward their ideal points and, if animus is sufficiently severe and symmetric, beyond. Meanwhile, when animus is high and asymmetric, we see party allegiances hold even as the substantive representation of one group's policy views altogether evaporates. As a result, rising levels of animus strengthen the appeal of parties that are either ideologically extreme or indistinguishable from one another.

The New Media Landscape. Much has been written about the downsides of an increasingly

fragmented media market in which consumers self-select into self-affirming news environments (for reviews, see Prior (2013); Winneg et al. (2014); Barberá (2020)). In this new media landscape, it is argued, communication is channelled through enclaves of like-minded individuals. Opportunities for persuasion, mutual understanding, and even the shared recognition of common facts all run in short supply. The result, say some, is a polity that is increasingly divided and "a breeding ground for extremism" (Sunstein, 2018, 71).[23]

Our model highlights yet another pathology associated with the proliferation of these exclusionary channels of political communication. Recall how candidates behave when given the opportunity to manipulate inter-group tensions. When their actions affect voters' understandings of both groups, as assuredly occurs when they are viewed by an entire electorate, candidates seek to temper animosities. By increasing the animus felt by all voters, after all, candidates are driven to the extremes of the policy spectrum without recovering any clear electoral reward. But when candidates can target their messages to voters more closely associated with their own group, incentives shift. Here, a candidate is more likely to win election when voters' enmity toward a group associated with the opposition grows. Rather than counsel inter-group tolerance and understanding, therefore, candidates in this setting stoke voters' animosities toward an out-group.

It is clear to see how the new media landscape puts us squarely in this latter world. But notice the source of the polarization that arises. The problem is not just that voters only hear views from political elites with whom they are aligned. The problem also is that voters are not privy to other communications in which their group is the target of hostilities. This structural asymmetry, in which voters rarely hear what the opposition is saying about groups to which they are affiliated, encourages political elites to foment inter-group hatreds. Changes in the media environment, as such, do not just facilitate new patterns of communication and behavior, as documented by the existing empirical literatures on the media. These changes also alter the political incentives for political elites to stoke animosities within an electorate.

[23] But see Barberá (2020) for a discussion of empirical studies that suggest that cross-cutting interactions occur more frequently than is usually supposed.

Separation of Powers. Our model reveals an interesting, under-appreciated, and vaguely ironic pathology of systems of separated power, which deliberately impede the translation of campaign promises into established law. In the United States, the Constitution's framers divided power among the various branches of government not just to guard against the accumulation of authority in any single individual or faction. They did so, also, as a check against the worst impulses of what they considered to be an ill-informed, unreasoning electorate. Through staggered elections and the distribution of state power across multiple and sometimes overlapping jurisdictions, it was thought, the government—and by extension, the nation itself—would be afforded some measure of protection against the turbulences and follies of popular sentiment.

James Madison makes the point most forcefully in *Federalist Paper 10*. In it, he recognizes the unavoidable "propensity of mankind to fall into mutual animosities" that yield factions intending to "vex and oppress each other." Unable to extinguish the "impulse of passion," the government instead must seek to control its effects. Because "the *causes* of faction cannot be removed," Madison insists, "relief is only to be sought in the means of controlling its *effects*." And the primary way of doing that is through the separation of powers.

At one level, Madison's design plainly succeeds. Separation of powers makes it nearly impossible for any one individual, no matter where she resides within the federal government, to advance a policy agenda all on her own. Consequentially, the outcome of any single election bears only weakly on the production of public policies, and the damage wrought by a would-be demagogue is mitigated.

Our model, however, reveals two other ways in which separation of powers affects the very political forces that Madison sought to contain. To begin, notice the effect that separation of powers has on candidate position-taking. Precisely because the probability of implementation is relatively small in systems of separated powers, voters incur a relatively low cost of insincere posturing (i.e., a low β in the language of our model). Consequentially, candidates have greater incentives to indulge the voter's animus—the baser passions and animosities, that is, that Madison lamented. So doing, the candidates assume increasingly extreme policy positions.

A second effect arises when animus is endogenized, and this one offers some reassurance for Madison's institutional design. As Proposition 3 tells us, the propensity of candidates to stoke inter-group rivalries decreases in the probability that a proposal becomes law (again, so that β is small). To the extent that separation of powers impedes lawmaking, therefore, it tempers the production of incendiary public appeals.

Separation of powers, we now can see, functions at different registers. In ways Madison plainly intended, this institutional design impedes the ambitions of a demagogue. But in ways he obviously did not, this same institutional design increases the chances that politicians will surrender to an electorate's appetite for demagoguery itself by assuming increasingly extreme policy positions.

7 Conclusion

We study a model of electoral competition in which candidates' positions are informed by voters' self-understandings and group enmities. By incorporating two expressive considerations into the voter's utility, we recover reasonably clear comparative statics on candidate position-taking, electoral fortunes, and the manipulation of inter-group animosities. All of these effects derive from changing electoral incentives for policy-motivated candidates.

Concerns about cognitive dissonance, we find, routinely discourage candidate extremism. The effects of animus, however, are more varied and subtle. When baseline levels of animus are reasonably low, marginal increases yield higher levels of elite polarization. At a higher baseline, meanwhile, increases in animus can push candidates to positions that are even more extreme than their own ideal points. And when this animus also is one-sided, we see both candidates attending to the policy preferences of the singularly aggrieved population. Consensus, as such, is not born of mutual accommodation or moderation, but rather of one candidate's capitulation to an enraged constituency with whom she has no affiliation.

The model also speaks to the strategic production of inter-group animosities. When candidates

can target their communications through siloed media channels, they have individual incentives to inflame inter-group tensions, even though doing so makes them collectively worse off. And this effect is particularly large when the stakes of an election are high, either because significant (un-modeled) power is vested in an office or because other (again un-modeled) factors ensure that campaign promises will translate into legislative achievements. Our model, as such, provides a counterpoint to the old-adage commonly attributed to Charles Philip Issawi, that "in any dispute the intensity of feeling is inversely proportional to the value of the issues at stake."

References

Abramowitz, Alan I and Steven Webster. 2016. "The Rise of Negative Partisanship and the Nationalization of U.S. Elections in the 21st Century." *Electoral Studies* 41:12–22.

Achen, Christopher and Larry Bartels. 2017. *Democracy for Realists*. Princeton, NJ: Princeton University Press.

Adams, James, Jay Dow and Samuel Merrill. 2006. "The Political Consequences of Alienation-based and Indifference-based Voter Abstention: Applications to Presidential Elections." *Political Behavior* 28(1):65–86.

Adams, James and Samuel Merrill. 2003. "Voter Turnout and Candidate Strategies in American Elections." *Journal of Politics* 65:161–189.

Ansolabehere, Stephen, Jonathan Rodden and James Snyder. 2006. "Purple America." *Journal of Economic Perspectives* 20(2):97–118.

Ash, Elliott, Massimo Morelli and Richard Van Weelden. 2017. "Elections and Divisiveness: Theory and Evidence." *The Journal of Politics* 79(4):1268–1285.

Barberá, Pablo. 2020. Social Media, Echo Chambers, and Political Polarization. In *Social Media and Democracy: The State of the Field, Prospects for Reform*, ed. Nathaniel Persily. Cambridge University Press pp. 33–55.

Baron, David P. 1994. "Electoral Competition with Informed and Uninformed Voters." *American Political Science Review* 88:33–47.

Beasley, Ryan K and Mark R Joslyn. 2001. "Cognitive Dissonance and Post-decision Attitude Change in Six Presidential Elections." *Political psychology* 22(3):521–540.

Boxwell, Levi, Matthew Gentzkow and Jesse Shapiro. 2024. "Cross-Country Trends in Affective Polarization." *Review of Economics and Statistics* 106(2):557–565.

Brennan, Geoffrey. 2008. "Psychological Dimensions in Voter Choice." *Public Choice* 137(3-4):475–489.

Brennan, Geoffrey and Alan Hamlin. 1998. "Expressive Voting and Electoral Equilibrium." *Public Choice* 95(1):149–175.

Callander, Steven and Catherine H Wilson. 2006. "Context-dependent Voting." *Quarterly Journal of Political Science* 1(3):227–255.

Callander, Steven and Catherine H Wilson. 2008. "Context-dependent Voting and Political Ambiguity." *Journal of Public Economics* 92(3-4):565–581.

Callander, Steven and Juan Carlos Carbajal. 2022. "Cause and Effect in Political Polarization: A Dynamic Analysis." *Journal of Political Economy* 130(4):825–880.

Clark, JR and Dwight R Lee. 2016. "Higher Costs Appeal to Voters: Implications of Expressive Voting." *Public Choice* 167(1):37–45.

Davis, Nicholas T. 2019. "Identity Sorting and Political Compromise." *American Politics Research* 47(2):391–414.

Ensley, Michael J. 2009. "Individual Campaign Contributions and Candidate Ideology." *Public Choice* 138(1-2):221–238.

Erikson, Robert and David Romero. 1990. "Candidate Equilibrium and the Behavioral Model of the Vote." *American Political Science Review* 84:1103–1126.

Fiorina, Morris P. 1976. "The Voting Decision: Instrumental and Expressive Aspects." *The Journal of Politics* 38(2):390–413.

Fiorina, Morris P, Samuel A Abrams and Jeremy C Pope. 2008. "Polarization in the American public: Misconceptions and misreadings." *The Journal of Politics* 70(2):556–560.

Fowler, Anthony, Seth J Hill, Jeffrey B Lewis, Chris Tausanovitch, Lynn Vavreck and Christopher Warshaw. 2023. "Moderates." *American Political Science Review* 117(2):643–660.

Glazer, Amihai. 1992. "An Expressive Voting Theory of Strikes." *Economic Inquiry* 30(4):733–741.

Glazer, Amihai. 2008. "Voting to Anger and to Please Others." *Public Choice* 134(3-4):247–254.

Hamlin, Alan and Colin Jennings. 2011. "Expressive Political Behaviour: Foundations, Scope and Implications." *British Journal of Political Science* 41:645–670.

Hamlin, Alan and Colin Jennings. 2018. Expressive Voting. In *The Oxford Handbook of Public Choice, Volume 1*, ed. Roger Congleton, Bernard Grofman and Stefan Voigt. New York, NY: Oxford University Press pp. 333–350.

Harmon-Jones, Eddie Ed and Judson Ed Mills. 1999. Cognitive dissonance: Progress on a pivotal theory in social psychology. In *Scientific Conferences Program, 1997, U Texas, Arlington, TX, US*. American Psychological Association.

Hill, Seth J and Chris Tausanovitch. 2015. "A Disconnect in Representation? Comparison of Trends in Congressional and Public Polarization." *The Journal of Politics* 77(4):1058–1075.

Hill, Seth J and Chris Tausanovitch. 2018. "Southern Realignment, Party Sorting, and the Polarization of American Primary Electorates, 1958–2012." *Public Choice* 176(1):107–132.

Hinich, Melvin J, John O Ledyard and Peter C Ordeshook. 1972. "Nonvoting and the Existence of Equilibrium under Majority Rule." *Journal of Economic Theory* 4(2):144–153.

Hirano, Shigeo, James M Snyder Jr, Stephen Daniel Ansolabehere and John Mark Hansen. 2010. "Primary Elections and Partisan Polarization in the US Congress." *Quarterly Journal of Political Science* 5:169–191.

Howell, William G and Terry M Moe. 2020. *Presidents, Populism, and the Crisis of Democracy*. University of Chicago Press.

Iyengar, Shanto, Gaurav Sood and Yphtach Lelkes. 2012. "Affect, not Ideology: A Social Identity Perspective on Polarization." *Public Opinion Quarterly* 76(3):405–431.

Iyengar, Shanto and Sean J Westwood. 2015. "Fear and Loathing across Party Lines: New Evidence on Group Polarization." *American Journal of Political Science* 59(3):690–707.

Iyengar, Shanto, Yphtach Lelkes, Matthew Levendusky, Neil Malhotra and Sean J Westwood. 2019. "The Origins and Consequences of Affective Polarization in the United States." *Annual Review of Political Science* 22:129–146.

Jones, Philip and John Hudson. 2000. "Civic Duty and Expressive Voting: Is Virtue its Own Reward?" *Kyklos* 53(1):3–16.

Kamenica, Emir and Louisa Egan Brad. 2014. "Voters, Dictators, and Peons: Expressive Voting and Pivotality." *Public Choice* 159(1):159–176.

Krasa, Stefan and Mattias K Polborn. 2018. "Political Competition in Legislative Elections." *American Political Science Review* 112(4):809–825.

Layman, Geoffrey C, Thomas M Carsey, John C Green, Richard Herrera and Rosalyn Cooperman. 2010. "Activists and Conflict Extension in American Party Politics." *American Political Science Review* 104(2):324–346.

Lee, Frances E. 2009. *Beyond Ideology: Politics, Principles, and Partisanship in the US Senate*. Chicago, IL: University of Chicago Press.

Levendusky, Matthew. 2009. *The Partisan Sort: How Liberals Became Democrats and Conservatives Became Republicans*. University of Chicago Press.

Lynam, Donald R and Joshua D Miller. 2019. "The Basic Trait of Antagonism: An Unfortunately Underappreciated Construct." *Journal of Research in Personality* 81:118–126.

Mason, Lilliana. 2015. ""I Disrespectfully Agree": The Differential Effects of Partisan Sorting on Social and Issue Polarization." *American Journal of Political Science* 59(1):128–145.

Mason, Lilliana. 2018*a*. "Ideologues without Issues: The Polarizing Consequences of Ideological Identities." *Public Opinion Quarterly* 82(S1):866–887.

Mason, Lilliana. 2018*b*. *Uncivil Agreement: How Politics Became our Identity*. Chicago, IL: University of Chicago Press.

May, John. 1973. "Opinion Structure of Political Parties: The Special Law of Curvilinear Disparity." *Political Studies* 21(2):135–151.

McCarty, Nolan. 2019. *Polarization: What Everyone Needs to Know*. New York, NY: Oxford University Press.

McCarty, Nolan, Keith T Poole and Howard Rosenthal. 2016. *Polarized America: The Dance of Ideology and Unequal Riches*. Cambridge, MA: MIT Press.

Moon, Woojin. 2004. "Party Activists, Campaign Resources and Candidate Position Taking: Theory, Tests and Applications." *British Journal of Political Science* 34(4):611–633.

Pietraszewski, David, Leda Cosmides and John Tooby. 2014. "The Content of Our Cooperation, Not the Color of Our Skin: An Alliance Detection System Regulates Categorization by Coalition and Race, but Not Sex." *PLoS ONE* 9(2).

Pietraszewski, David, Oliver Curry, Michael Petersen, Leda Cosmides and John Tooby. 2015. "Constituents of Political Cognition: Race, Party Politics, and the Alliance Detection System." *Cognition* 140:24–39.

Polborn, Mattias K and James M Snyder Jr. 2017. "Party Polarization in Legislatures with Office-motivated Candidates." *The Quarterly Journal of Economics* 132(3):1509–1550.

Prior, Markus. 2013. "Media and Political Polarization." *Annual Review of Political Science* 16:101–127.

Riker, William H and Peter C Ordeshook. 1968. "A Theory of the Calculus of Voting." *The American political science review* 62(1):25–42.

Schuessler, Alexander A. 2000. "Expressive Voting." *Rationality and Society* 12(1):87–119.

Sides, John, Michael Tesler and Lynn Vavreck. 2018. *Identity Crisis: The 2016 Presidential Campaign and the Meaning of America*. Princeton, NJ: Princeton University Press.

Sunstein, Cass R. 2018. *Republic.com*. Princeton, NJ: Princeton University Press.

Tajfel, Henri. 1970. "Experiments in Intergroup Discrimination." *Scientific American* 223(5):96–103.

Taylor, Brad R. 2015. "Strategic and Expressive Voting." *Constitutional Political Economy* 26(2):159–170.

Theriault, Sean. 2008. *Party Polarization in Congress*. New York, NY: Cambridge University Press.

Valentino, Nicholas A, Ted Brader, Eric W Groenendyk, Krysha Gregorowicz and Vincent L Hutchings. 2011. "Election night's alright for fighting: The role of emotions in political participation." *The Journal of Politics* 73(1):156–170.

Webster, Steven. 2020. *American Rage: How Anger Shapes our Politics*. Cambridge Univ Press.

Winneg, Kenneth M, Daniel M Butler, Saar Golde, Darwin W Miller III and Norman H Nie. 2014. Online News Consumption in the United States and Ideological Extremism. In *The Oxford Handbook of Political Communication*. pp. 809–822.

Zmigrod, Leor. 2020. "The Role of Cognitive Rigidity in Political Ideologies: Theory, Evidence, and Future Directions." *Current Opinion in Behavioral Sciences* 34:34–39.

8 Online Appendix

Proof of Lemma 1. In case that $\mu_j < x_L, x_R$ or $\mu_j > x_L, x_R$, simplifying (2) yields

$$\beta|x_R - \bar{\theta}_i| - \beta|x_L - \bar{\theta}_i| = \begin{cases} \alpha(x_R - x_L) & \text{if } \mu_j < x_L, x_R \\ -\alpha(x_R - x_L) & \text{if } \mu_j > x_L, x_R \end{cases} \tag{20}$$

If, on the left-hand side, the two terms inside the absolute values are either both positive or both negative, the equation simplifies to $\alpha = \beta$ (i.e., the case that all individuals are indifferent between all policy positions), or $\alpha = -\beta$, a contradiction since both parameters are positive by assumption.

Thus, consider the case $x_L < \bar{\theta}_i < x_R$. Simplifying (2) then yields the expressions in the statement of the lemma.

In case that $x_L < \mu_j < x_R$, attempting $\bar{\theta}_i < x_L, x_R$ or $> x_L, x_R$ leads to an immediate contradiction. Thus, conjecture that $\bar{\theta}_i \in [x_L, x_R]$. Then, simplifying (2) yields the expression in the statement of the lemma. ■

Proof of Lemma 2. Recall that the distributions Φ_i are symmetric around zero. Thus, if $q_1 = q_2$ and $\Phi_1 = \Phi_2$ then (4) holds if and only if $\bar{\theta}_1 - m - \mu_1 = -(\bar{\theta}_2 - m - \mu_2)$. Solving this equation for m yields (5). ■

Lemma 3 *Any solution to the first-order conditions* (10) *or* (11) *solves the candidates' optimization problems.*

Proof of Lemma 3. Divide the left-hand side of the first-order condition (10) by $F(m)$ and rearrange terms to get

$$\frac{f(m(x_L, x_R))}{F(m(x_L, x_R))}\left(\frac{\beta - \bar{\alpha}}{2\beta}\right)(x_R - x_L) - 1 = 0. \tag{21}$$

Symmetry of F implies that $F(m) = 1 - F(-m)$ and $f(m) = f(-m)$. Thus, (21) is equivalent to

$$\frac{f(-m(x_L, x_R))}{1 - F(-m(x_L, x_R))}\left(\frac{\beta - \bar{\alpha}}{2\beta}\right)(x_R - x_L) - 1 = 0. \tag{22}$$

Equation (6) implies that $m(x_L, x_R)$ is increasing in x_L. Hence, $-m(x_L, x_R)$ is decreasing in x_L. Because the distribution, F, has an increasing hazard rate, and $-m$ decreases, the first factor in (22) is decreasing. Furthermore, $x_R - x_L > 0$ and decreases in x_L. Thus, (21) is strictly decreasing in x_L, which implies that the solution to the first-order condition is a maximum.

The argument for the first-order condition with respect to x_R is similar and omitted. ■

Proof of Proposition 1. The first two statements are in the text. To prove the last statement, differentiate (13) with respect to β to get

$$\frac{\partial x_R}{\partial \beta} = \frac{[f(0)\,(2\beta - 2\bar{\alpha})] - 2f(0)[\beta + f(0)\mu(\alpha_2 - \alpha_1)]}{[f(0)\,(2\beta - \alpha_1 - \alpha_2)]^2} = \frac{-2f(0)[\bar{\alpha} + f(0)\mu\Delta\alpha]}{[2f(0)\,(\beta - \bar{\alpha})]^2}, \tag{23}$$

which is always negative. Thus, an increase in β leads to a leftward shift in x_R. Similarly,

$$\frac{\partial x_L}{\partial \beta} = -\frac{f(0)\,(2\beta + \Delta\alpha) - [\beta - f(0)\mu\Delta\alpha]2f(0)}{[f(0)\,(2\beta - \alpha_1 - \alpha_2)]^2} = \frac{2f(0)[\bar{\alpha} - f(0)\Delta\alpha]}{[2f(0)\,(\beta - \bar{\alpha})]^2}. \tag{24}$$

This is positive if and only if the term in square brackets is positive; thus, in particular, if $\Delta\alpha$ is sufficiently close to zero.

Finally, we show that $x_R > \mu$. Substituting $1/(2\mu)$ for $f(0)$ into (13) yields $x_R = ((\alpha_1 - \alpha_2 - 2\beta)\mu)/(\alpha_1 + \alpha_2 - 2\beta) > \mu$ as $\alpha_i < \bar{\beta}$. Finally, (13) is strictly decreasing in $f(0)$. Thus, x_R is also strictly larger than μ for $f(0) < 1/(2\mu)$. The argument for $x_L < \mu_L$ is analogous. ■

Proof of Proposition 2. If $x_R < \theta_R$, then Candidate R's first order condition is

$$-f(m(x_L, x_R))\left(\frac{1}{2} - \frac{\Delta\alpha}{4\beta}\right)(x_R - x_L) + (1 - F(m(x_L, x_R))) = 0. \tag{25}$$

Note that $x_L < \theta_R$, else (16) cannot hold. If $\Delta\alpha \to 2\beta$ then $m \to x_L$, which implies that $F(m) \leq F(\theta_R) < 1$. However, as $\Delta\alpha \to 2\beta$ the first term in (25) converges to zero, and hence $F(m)$ must converge to 1, a contradiction. Thus, the derivative of candidate R's objective function, i.e., the left-hand side of first-order condition (25), is strictly positive at $x_R = \theta_R$ for $\Delta\alpha$ close to 2β.

Next, it is not possible for $x_R > \theta_R$ as long as $\Delta\alpha < 2\beta$. In particular, if $x_R > \theta_R$ the derivative of candidate R's objective is

$$-f(m(x_L, x_R))\left(\frac{1}{2} - \frac{\Delta\alpha}{4\beta}\right)(2\theta_R - (x_R + x_L)) - (1 - F(m(x_L, x_R))) = 0, \tag{26}$$

Note that $2\theta_R > (x_R + x_L)$. In particular, if $\theta_R \leq x_L < x_R$ then Candidate R can improve by deviating to $x_R = x_L$ because the winning position is closer to candidate R's ideal point. Further, it is always optimal for Candidate L to choose a position that is strictly to the left of x_R. Hence $x_L < x_R \leq \theta_R$, which implies that $2\theta_R > (x_R + x_L)$. The left-hand side of (26) is therefore strictly negative, which implies that it is optimal for Candidate R to choose $x_R = \theta_R$.

At $x_R = \theta_R$ we therefore get

$$-f(m(x_L, \theta_R))\left(\frac{1}{2} - \frac{\Delta\alpha}{4\beta}\right)(\theta_R - x_L) + (1 - F(m(x_L, \theta_R))) \geq 0. \tag{27}$$

Adding this to Candidate L's first-order condition (16) yields

$$f(m(x_L, \theta_R))\frac{\Delta\alpha}{2\beta}(\theta_R - x_L) + 1 - 2F(m(x_L, \theta_R)) \geq 0. \tag{28}$$

Thus, $1 - 2F(m) < 0$ and hence $F(m) > 0.5$, i.e., candidate L wins with a probability that is strictly greater than 0.5.

It remains to prove that Candidate L's second-order condition is satisfied. Note that symmetry of the distribution implies $F(m) = 1 - F(-m)$ and $f(m) = f(-m)$. Divide both sides of (16) by

$1 - F(-m)$ to get

$$\frac{f(-m(x_L, \theta_R))}{1 - F(-m(x_L, \theta_R)} \left(\frac{1}{2} + \frac{\Delta\alpha}{4\beta} \right) (x_R - x_L) - 1 = 0 \tag{29}$$

m is increasing in x_L, and hence $-m$ is decreasing in x_L. Thus, the hazard rate $f(-m)/(1 - F(-m))$ is decreasing in m. Further $(x_R - x_L)$ decreases as x_L is increased. Hence the left-hand side of (29) is strictly decreasing in x_L at any solution of the first-order condition. This, in turn, implies that the left-hand side of (16) is strictly positive marginally to the left of x_L and strictly negative marginally to the right of x_L. Hence candidate L's optimization problem does not have a local minimum. This implies that the solution of (16) is both unique and a maximum.

Finally, we want to prove that, when $x_R = \theta_R$, then an increase in $\Delta\alpha$ does not affect the winning probability. Solving (15) for the value of x_L that keeps m constant yields

$$\tilde{x_L} = \frac{4\beta m - (2\beta - \Delta\alpha)\theta_R}{2\beta + \Delta\alpha}$$

Substituting $\tilde{x_L}$ in the first-order condition (16) and simplifying yields

$$f(m)(\theta_R - m) - F(m) = 0,$$

which is independent of $\Delta\alpha$; therefore, marginal changes in $\Delta\alpha$ leave the equilibrium m, and thus the winning probability, unaffected. ■

Proof of Proposition 3.

Let $\tilde{\alpha}$ be the value at which (12) and (13) are equal to θ_L and θ_R, respectively for $\alpha_1 = \alpha_2 = \tilde{\alpha}$. Thus, $\tilde{\alpha}$ solves

$$\frac{\hat{\beta}}{2f(0)(\hat{\beta} - \tilde{\alpha})} = \theta_R. \tag{30}$$

Hence, the value of $\tilde{\alpha}$ is the one given in the proposition. Note $\theta_L < x_L < x_R < \theta_R$ when $\alpha_1 = \alpha_2 < \tilde{\alpha}$. In the main text we have shown that an equilibrium with $\alpha_i > 0$ and $\theta_L < x_L < x_R < \theta_R$ cannot exist. Thus, in a symmetric equilibrium either $\alpha_1 = \alpha_2 = 0$ or $\alpha_i \geq \tilde{\alpha}$, for $i = 1, 2$. This proves the

first statement.

Suppose that $\alpha_1 = \alpha_2 = 0$. Because of convexity of payoffs in α_i when $\theta_L < x_L < x_R < \theta_R$, it is sufficient to consider deviations to and α_i at which candidate i's policy just moves to the ideal point θ_i, $i = L, R$. Further, without loss of generality we can restrict attention to deviations by candidate R. Also, note that

Let $\alpha_1 = \alpha_2 = 0$. Then substituting (10) and the fact that the winning probabilities are 0.5 into (18) implies that candidate R's utility is

$$-p\frac{1}{2}\left(\theta_R - \frac{1}{2f(0)}\right) - p\frac{1}{2}\left(\theta_R + \frac{1}{2f(0)}\right) = -p\theta_R. \tag{31}$$

If $\alpha_1 = 0$ then the level of animus that candidate R would need to reach $x_R = \theta_R$ is given by

$$\tilde{\alpha}_2 = \hat{\beta}\frac{2\theta_R f(0) - 1}{f(0)(2 + \theta_R)}. \tag{32}$$

Substituting this value of $\tilde{\alpha}_2$ into (10) implies that x_L does not depend on $\hat{\beta}$. Because the winning probabilities are 0.5, (18) implies that candidate R's payoff after a deviation to $\tilde{\alpha}_2$ equals

$$-\frac{1}{2}p(\theta_R - x_L) - b\tilde{\alpha}_2. \tag{33}$$

In order for no-animus to be an equilibrium, the payoff from deviation must be less or equal to the equilibrium payoff (31), which is equivalent to

$$\frac{1}{2}(\theta_R - x_L) + \frac{b\tilde{\alpha}_2}{p} \geq \theta_R. \tag{34}$$

Note that $\tilde{\alpha}_2$ is increasing in $\hat{\beta}$. Thus, (34) is easier to satisfy if $\hat{\beta}$ is increased of p is decreased.

Let $\bar{b}$ be the value of b at which (34) holds with equality. Then

$$\bar{b} = \frac{3p(\theta_R - x_L)}{2\tilde{\alpha}_2}. \tag{35}$$

Thus, $\bar{b}$ is increasing in p and decreasing in $\hat{\beta}$ (through $\tilde{\alpha}_2$), which proves the second statement. ■