Policy Anchors, Voter Behavior, and the Dynamics of Party Platforms*

Peter Bils† Steven Callander‡ Federica Izzo§

Abstract

Voting is a complex task, and even well-meaning voters must use shortcuts and heuristics to make decisions. In this paper we develop a model in which voters use past platforms as reference points when evaluating parties' current proposals. We show that these *policy anchors* lead to preference reversals in voting—both within and across elections—relative to the standard model. We then incorporate policy anchors into a canonical model of electoral competition and show that anchors change policy outcomes substantively. Party platforms respond to anchors non-monotonically, such that centrist and extreme anchors both induce extreme platforms. This causes policy platforms to evolve continuously where they would otherwise be largely stable, and for platform cycles to necessarily be asymmetric, with large jumps outward followed by incremental movements inward. We connect these results to broader questions about voter welfare and political representation, as well as observed patterns in platform dynamics.

*For helpful comments, we thank Neil Malhotra, Peter Buisseret, Juan Carlos Carbajal, Gilat Levy, Helios Herrera, and seminar audiences at APEN, IE-Rochester, Princeton-Warwick-Yale, SPSA, and Vanderbilt.

†Vanderbilt University, Departments of Political Science and Economics; peter.h.bils@vanderbilt.edu.

‡Stanford University, Stanford Graduate School of Business; sjc@stanford.edu.

§Vanderbilt University, Departments of Political Science and Economics; federica.izzo@vanderbilt.edu.

1 Introduction

A half century of research into political behavior has established all too clearly that voters fall short of any notion of a democratic ideal. Voters are neither omniscient information processors nor perfectly rational decision makers.

It does not follow from these shortcomings, however, that voters are irrational or, worse still, indifferent to political outcomes. Evaluating parties' policy programs is a difficult problem for voters, and given the cognitive limitations of the human mind, it is to be expected that even well-meaning decision makers must adopt a myriad of shortcuts and rules-of-thumb in the pursuit of effective decision making (Simon, 1955).

One natural way in which voters may attempt to cut through this complexity involves using the past as a guide to evaluate current alternatives. Studies from behavioral economics have repeatedly documented that people's evaluation of their current alternatives is formed by making comparisons to the past, even when the past should be irrelevant to the decision at hand.[1] We propose that the vote choice is no different. That voters make sense of politics by using the past as a guide. Specifically, that voters evaluate and respond to policy platforms based, in part, upon positions that have been offered by parties in earlier elections. A voter sees a party's platform as more attractive if it is closer to her ideal than was the party's platform in the last election.

To provide a concrete example, the "compassionate conservatism" of George W. Bush was attractive to voters as it provided a positive contrast to the sharper-edged conservatism of an earlier era. The new policy offering can be seen by the voter as more moderate because relative to past policies it is. In this way, past policy positions of the parties serve as *anchors* for the voting decision.

In this paper, we formalize a theory of voting with policy anchors. The use of anchors distorts voter behavior relative to the standard model of spatial voting. Voters may support a party proposing a platform farther from their ideal, and reverse their choice across elections even if faced with identical alternatives. Incorporating policy anchors into a classic model of two-party competition, we show they produce rich dynamics in party platforms. Within each election, party platforms respond to anchors non-monotonically. Moderate anchors induce moderate platforms, whereas centrist and extreme anchors both induce more extreme platforms. Across elections, this U-shaped reaction function causes party platforms to cycle continuously between moderate and more extreme positions in an environment where change would otherwise be infrequent. Furthermore, such cycles are *necessarily* asymmetric, with

[1] For example, people make decisions based on how their current option compares to their previous choice sets (Tversky and Simonson, 1993), their status quo endowment (Kahneman and Tversky, 1979), and unrelated information (Tversky and Kahneman, 1974).

incremental movements inward punctuated by large jumps outward.

Our findings match the frequent movements, both small and large, of party platforms in elections across a wide variety of countries and contexts. Moreover, our model provides a rationale for the striking finding in the empirical literature that platform dynamics tend to exhibit larger variance than trends in voters' underlying policy opinions (Page and Shapiro, 1992; Budge et al., 2010; Druckman and Leeper, 2012).[2]

Policy Anchors and Constructed Preference

A long-standing perspective in the psychology of choice is that preference is constructed rather than inherent (Lichtenstein and Slovic, 2006). A key insight in this literature is that "... the true objects of evaluation and choice are neither objects in the real world nor verbal descriptions; they are mental representations" (Tversky and Kahneman, 2000, p. xiv). The idea, then, is that when individuals must evaluate an alternative, they rely on comparable "reference points" that lie outside of their choice set but offer a useful benchmark for calibration (Kőszegi and Rabin, 2006). Because such reference points may be context-dependent, or evolve over time, people lack stable consistent preferences over alternatives. Instead, they construct their preference anew for each decision problem.

Voting is a natural domain for preference to be constructed rather than inherent. The voting decision is extraordinarily broad, covering many distinct issue domains, with information that is dispersed and often hard to integrate. Politics is difficult to comprehend, even for experts. Making voting even more difficult is that the policy space is *scale-free*. In consumer markets, for example, products typically have concrete uses, the efficiency and attractiveness of which can be directly compared to other products, and they come with a price that provides a metric for comparison. In politics, in contrast, policy positions stand alone, with no natural "score" or metric for evaluation, making comparisons difficult, especially for policies on opposite sides of the ideological spectrum. What does it mean to an average voter, for example, if Republicans offer a tight immigration policy? How much should the voter value this policy? In particular, how does she compare this proposal, that may lie to her right, to the Democrats' position that is on her left? In this domain, then, it seems natural that voters may be constructing "mental representations" (Tversky and Kahneman, 2000, p. xiv) of the different policy proposals when choosing which party to support, and that they may use reference points to do so.

Our theory of voting with anchors builds on these ideas in two key ways. First, voters use past platform offerings as reference points, to anchor their evaluation of today's alternatives. The large literature analyzing reference points in both psychology and economics shows that

[2]The extent of voter movement is an ongoing debate in the academic literature (e.g., Boxell et al., 2024), although a consensus does hold that voters have shifted ideological ground *less* than have elites.

individuals evaluate alternatives in terms of gains or losses relative to the reference point, whether that is their status quo utility, their expectations, or some aspiration level (Kahneman and Tversky, 1979; Kőszegi and Rabin, 2006). In our model, it is the parties' previous platform offerings that offer this benchmark for evaluation. This is in line with the literature on choice and background context (Tversky and Simonson, 1993), which shows that the previous choice set influences an individual's evaluations of her current alternatives. Thus, voters in our model evaluate today's proposals more positively if they compare favorably to past platforms. More precisely, a platform will look more appealing to the voter if, compared to previous offerings, it represents a shift closer to her ideal. Conversely, if the platform is seen by the voter as a "loss" relative to the reference point, this worsens her evaluation.

Second, we assume that these comparisons are directional and side-specific. Voters construct their preferences by comparing a policy to past platforms on the same side of the spectrum. A voter sees a policy today that is to the left of her ideal point as more (less) appealing if it is relatively closer to (farther from) her than the previous left-wing platform. Likewise, her evaluation of policies to the right of her ideal point is anchored to past right-wing offerings. Our approach is grounded in the idea that, faced with a scale-free policy space, a voter will have difficulty directly comparing policies to the left of her ideal point to policies to the right of her ideal point. This implies that, if voters use reference points to construct preferences, those reference points must be side-specific. If voters could readily compare a left-wing policy today to a right-wing policy yesterday, then they could just as easily compare today's cross-spectrum options and would not need anchors.

This assumption on the side-specific effect of comparisons across platforms aligns with evidence from the political behavior literature. Using a survey experiment based on US elections, Simonovits (2017) demonstrates how the addition of an extreme candidate on one side of the spectrum makes a moderate candidate from the same side more appealing. In fact, when asked to place the moderate conservative on the policy spectrum, participants who were first exposed to an extreme conservative placed a moderate candidate closer to the center.[3] Similarly, exposure to an extreme liberal candidate induced respondents to place the moderate liberal candidate closer to the center of the spectrum.[4] Consistent with our theory, voters, in effect, react to the presence of the extreme conservative by rescaling the right side of the policy space to fit him in, and to the extreme liberal by rescaling the left side. Moreover, in line with our assumption, this rescaling is directional in that a moderate liberal candidate is immune from comparison to an extreme conservative on the other side

[3] Rotter and Rotter (1966) provide similar experimental evidence, albeit with dated methods, on the 1964 U.S. presidential election examining the effect of Wallace on Goldwater's vote total.

[4] Callander and Wilson (2006) provide observational evidence that the local context affects the turnout decision of voters.

of the political landscape, and vice versa. Waismel-Manor and Simonovits (2017) and Wang and Chen (2019) provide related evidence from Israel and Taiwan, respectively.

We note that these studies examine settings with *multiple candidates*, and document the one-sided effect of comparisons across platforms on the same side of the spectrum within an election. We theorize that when a voter evaluates only two candidates—one from each side—past platform offerings play a similar reference-point role.[5]

More Detail on the Model and Our Contribution

Our contribution has two connected parts. First, we begin by introducing a formal model of voting with policy anchors. We show that policy anchors lead to two distinct types of preference reversals relative to the standard model of spatial voting. First, *within* an election, a voter using a policy anchor may vote for the party that is more distant from her in the policy space, even though, all else equal, she prefers policies that are closer to her ideal point. Second, *across* elections, a voter may switch her preference over the same pair of policy positions from one election to the next. Although both preference reversals suggest that standard spatial voting is violated, we show that both types of behavior are natural consequences of spatial voting when voters evaluate positions using policy anchors.

Our second contribution is to explore how policy anchors shape politics more broadly. We ask two questions: How do parties respond strategically to voters using policy anchors? And what does this mean for the dynamics of party platforms? We build a model of electoral competition to answer these questions. The model we build incorporates policy anchors into the classic model of two-party competition with uncertainty of Wittman (1983) and Calvert (1985) while extending it dynamically to an arbitrary number of elections.

Voter anchoring changes the incentives of the parties. Because a party will appeal more to voters if it moves toward the center from one election to the next, policy anchoring gives the parties an extra incentive to converge to the median voter. However, a key finding of the model is that, rather than moderating each period, policy platforms can also become more divergent, and in fact the jumps outward are generally larger than movements inward.

We begin our analysis by showing that policy platforms of the parties depend crucially on how voters use policy anchors. In line with the reference-points literature, the marginal effect of a platform moving away from the reference point for voters in our model is larger for smaller movements, and flattens as the platform is farther from the anchor. As a consequence, the most moderate party platforms emerge when the policy anchors are moderate, whereas

[5] Most relevant in the theoretical literature is Callander and Wilson (2006), which shows how the local voting context can shape the turnout decision *within* a single election but that it cannot affect the vote choice. In contrast, we show that the background context—*across* elections—can affect the vote decision itself, even with only two parties competing. Callander and Wilson (2008) extend the model to allow for ambiguity in party platforms.

centrist and extreme anchors both produce extreme platforms. Thus, the response of party platforms is not monotonic in the policy anchor, rather the platforms react in a U-shape.

Across elections, the U-shaped reaction function generates platform reversals in equilibrium: parties move to the extreme today, then reverse course and creep back toward the median voter tomorrow.

It may seem surprising that a party would ever shift its platform *away* from the median voter when voters are anchoring. A move away from the median makes a party *less* attractive than in the standard model without anchoring. The surprise, however, is not the outward move itself. Starting from a centrist position, a party gains a policy benefit from shifting outward—by the classic Calvert–Wittman logic—and this benefit can outweigh the anchoring cost in voter evaluations. Parties shift outward not because of anchoring, but despite it.

What is surprising is that this outward shift is then reversed. How can it be optimal to move away from the median in one election if the best response is to shift back in the next? The answer lies in the diminishing marginal effect of distance from the anchor: moving away has a larger effect when close to the anchor and a smaller effect for platforms already far from it. When the anchor is initially very moderate, the benefit from further convergence is small, so parties have an incentive to diverge, and—because the effect of anchoring is diminishing in distance—they do so in a large jump. But once they move, the anchor moves with them to the extreme location, restoring the benefit of moving back toward the median and inducing an inward shift to improve electoral performance. In a sense, parties overshoot in reaction to policy anchors by moving to the extreme, then undo that overreaction. The diminishing impact of anchoring creates, in effect, a time-inconsistency in party incentives, which drives reversals in their policy platforms.

This process can generate long-lasting dynamics even in an environment that is otherwise unchanging, as platforms jump, reverse course, and converge to a final resting place. This apparent long-run stability, furthermore, is illusory. In an environment that does itself change, subject to periodic shocks, anchoring leads to party platforms that cycle continuously, jumping outward and creeping inward before restarting the process once again, even if the shocks are small and infrequent and platform change would otherwise be rare without anchoring.

Underlying the effects of voter anchoring are deep questions about political representation, voter welfare, and measurement. Do the movements in party platforms make voters better off if voters' ideal policies are fixed and unchanging? Should we even measure voter welfare by these ideal points—what behavioral economists would refer to as voters' *true* or *fundamental* preferences—or should welfare include the anchoring effect, the *decision utility* voters actually experience when casting their votes?

This question surfaces a fundamental ambiguity in the notion of political representation.

Representative democracy requires that parties *represent* the preferences of voters. If we take voters' true preferences as the standard, then the constant shifting of party platforms in our model implies a failure of political representation. But as the party movements are in response to the voters' anchoring—their decision utility from voting—should we not view this as the pure instantiation of political representation? Our paper offers a behavioral perspective on these questions and whether policymakers act as delegates or trustees in representing voter interests (see, e.g., Pitkin, 1967). In emphasizing the importance of voters' cognitive limitations in understanding political representation, our theory contrasts with the standard perspective which emphasizes voters' lack of expertise. We return to this and other thorny questions after presenting the formal analysis.

It is important to note that our formalization is meant to *represent* voter preferences constructed in this way, rather than be a *literal* description of how a voter forms her preference. Following the method of mathematical psychology, it is not that the voter literally calculates her utility. Rather, it is that her behavior in practice can be represented, or described, by a particular utility function.

Other Related Ideas and Literature

As detailed above, our model builds on the reference-point literature originating with Kahneman and Tversky (1979). We depart in some important ways, which we now discuss. Most of that literature couples reference points with loss aversion, evaluating gains and losses asymmetrically. In our model, by contrast, platform shifts relative to the reference point—toward or away from the voter—are treated symmetrically. This choice avoids hard-coding asymmetries into the model and lets us isolate a distinct, strategic force that generates asymmetric platform dynamics in equilibrium.

It is also worth noting that, while much of this tradition focuses on experienced gains and losses and thus restricts reference points to concrete outcomes, the original formulation in Kahneman and Tversky (1979) is broader: reference points can reflect experienced utility, expectations, or aspirations. In our model, a previously proposed platform influences a voter's evaluation today even if it was never implemented.

Moreover, specific to our application to elections and differently from the economics literature, in our model the voters consider different reference points for different alternatives. Voters' use of policy anchors is side specific, as they assess each party's current platform relative to past positions on the same side of the ideological spectrum. This aspect also distinguishes our work from previous applications of reference points to politics, which all assume that voters use the same reference point (typically, the status quo) to evaluate all policies or candidates (see Alesina and Passarelli (2019); Lockwood and Rockey (2020); Grillo and Prato

(2023)). This difference is crucial to the U-shaped response of party platforms to the anchors, and thus to the cyclic dynamics that we uncover. Furthermore, these existing works focus on characterizing parties' behavior within a single period, in contrast with our objective of studying the long-run evolution of platforms.[6] Lockwood et al. (2024) analyzes long-run platform evolution, but again in a setting where the status-quo utility is the reference point that alters voters' evaluation of both platforms. They exploit loss aversion—voters' asymmetric reaction to gains and losses—to generate *monotonic* changes in platforms. In contrast, loss aversion plays no role in our mechanism. Our result is driven instead by the interaction of the anchor and party incentives, and we show this generates a *non-monotonic* response of party platforms to the anchor.

Beyond the literature on reference points, other works within the behavioral formal-theory tradition embed psychological insights into game-theoretic models of electoral competition. Several papers within this niche explore the dynamics of party platforms. Of these, closest to our model is Callander and Carbajal (2022). They study the comovement of party platforms and voter preferences and explain polarization but not cycles, whereas we fix voters' innate preferences and show how cycles can emerge. Levy and Razin (2025) provide a model of cycles based on voter learning about the best policy. Our model is one of complete information.

Our theory of policy anchors proposes one way in which voters might look to past policy platforms to evaluate a candidate. One very different way is through a reputational lens in which a candidate's true position is a weighted average of the positions he has held. Such a theory would lead to an opposite effect to what drives our theory, namely that a candidate's platform today is less appealing if his platform yesterday was worse. A reputational approach represents a distinct mechanism and a fundamentally different underlying psychological process to what we model, one that is best grounded in belief updating rather than the foundation of constructed preference that we build upon. The two approaches also differ in other ways. For instance, voter anchoring is not candidate specific as a reputational theory would be. Indeed, the comparisons that drive anchoring do not need to come from the same person or entity, and need only be in the same direction as the platform under consideration.[7] Thus, our theory allows us to gain insight into voters and parties' behavior even when candidates change across elections.

Finally, our results on the dynamics of policy platforms connect with the important liter-

[6] Grillo (2016) considers a game where candidates can make statements about their (initially unobserved) valence, and these statements form the reference points used by voters. As such, voters also have candidate-specific reference points. However, in Grillo (2016) there is only one election and reference points are over valence, rather than policy. As such, the paper studies a very different set of questions from our work.

[7] To make the distinction even clearer, our theory can also apply to voting on ballot measures without an individual or entity attached.

ature on path dependence (Pierson, 2000). Page (2006) points out that path dependence, as used in the literature, is not a single concept and can take several different forms. This is the case in our model as well. Following Page's (2006) terminology, the cycles we identify in policy platforms are *path*-dependent but are not *outcome*-dependent. This means that the outcome in any election depends on the history—specifically, on the policy platforms offered by the parties in the previous election—but that the distribution of outcomes over the long-run is independent of the history.

2 Policy Anchors and Voter Behavior

In this section, we introduce policy anchors to the spatial theory of voting. We then explore the implications of anchors for voting behavior.

2.1 A Theory of Voting with Anchors

We consider a standard spatial voting environment. Throughout, we take the policy space to be single dimensional and given by $[-1, 1]$, with a generic policy platform at time t given by p_t.

Voters possess ideal points in $[-1, 1]$ and evaluate policies spatially. For now, consider a single voter with an ideal policy of $x \in [-1, 1]$ who at time t is forming an opinion of a policy platform to the right of her ideal point, which we denote $r_t > x$. In the reference point literature, individuals' (perceived) utility from an alternative has two components: 'intrinsic utility', which corresponds to the standard 'outcome-based' utility, and 'gain-loss utility', which reflects how the alternative under evaluation compares to the reference point (Kőszegi and Rabin, 2006). Analogously, in our model the voter's evaluation depends on both the direct distance of the platform to her ideal policy (the intrinsic component) *and* how that platform looks relative to the previous policy offering to the right of her ideal point, $r_{t-1} > x$ (the gain-loss component). Likewise, when evaluating a policy to the left of her ideal point, denoted $d_t < x$, she contrasts it to the previous platform that was offered to her left, $d_{t-1} < x$. The policies r_{t-1} and d_{t-1} are the *anchors* that the voter uses to help her evaluate platforms and construct her preferences over policies.

The relative distance of the policy platform $p_t \in \{d_t, r_t\}$ is captured by the following:

$$\Delta p_t = |p_{t-1} - x| - |p_t - x|.$$

When the voter considers a platform and anchor to the right of her ideal policy, this relative distance is given by $\Delta r_t = r_{t-1} - r_t$, while if the platform is to the left of x it reduces to

$\Delta d_t = d_t - d_{t-1}$. Thus, a positive value of Δp_t means the policy is relatively more appealing to the voter than it was previously, and a negative value that it is less appealing.[8]

The voter's preference combines this intertemporal evaluation with a direct evaluation of the party, and can be represented by the following utility function:

$$u_x(p_t|p_{t-1}) = \underbrace{-|p_t - x|}_{\text{intrinsic}} \quad + \quad \underbrace{f(\Delta p_t)}_{\text{gain-loss}}, \tag{1}$$

The intrinsic component of utility follows preference in the standard way, with platforms closer to the voter's ideal policy receiving higher utility than those further away. We assume this component is linear in policy to highlight clearly the interaction with the indirect component.

How the gain-loss component impacts utility depends on the function, f. We assume f is twice-differentiable. Consistent with our theory, we take f to be strictly increasing in Δp_t, $f'(\Delta p_t) > 0$, so that relative movements away from the voter are viewed unfavorably. We also assume that the effect of the anchor is neutral in the following sense:

$$f(0) = 0, \quad \text{and} \quad f(a) = -f(-a). \tag{2}$$

Thus, if the party platform does not shift, the indirect effect is zero, and the boost from shifting toward the voter is exactly matched by an equal shift away. Our results do not rely on this symmetry. We impose it for clarity and to highlight that the asymmetry in gains versus losses that drives prospect theory is not necessary for our results.

The literature on reference points assumes that the marginal effect of the reference point on an individual's evaluation is stronger for small shifts away from the reference, and diminishes as the alternative becomes more distinct from it (see, e.g., Kahneman and Tversky (1979) for theoretical arguments on this point and Abdellaoui (2000) for empirical evidence). In other words, the effect is convex for alternatives that constitute a loss relative to the reference point, and concave for alternatives that constitute a gain. In terms of our set-up, this corresponds to an s-shaped f function: the marginal effect of today's platform moving away from the relevant anchor—whether toward the voter or farther from her—becomes smaller for larger movements. Formally, we assume $f''(\Delta p_t) > 0$ for $\Delta p_t < 0$ and $f''(\Delta p_t) < 0$ for $\Delta p_t > 0$, as depicted in Figure 1.

For voters on the flanks, it may be that both platforms from the previous period are on the same side of her ideal point. For example, suppose $d_{t-1} < r_{t-1} < x$. In this case, for a

[8] The dependence on Δp_t implies that the effect of a movement of certain size is independent of the distance the anchor is from the voter. This can be generalized to allow the absolute distance of the platforms from the voter to impact the evaluation through the use of *skew-symmetric* functions (Fishburn, 1982).

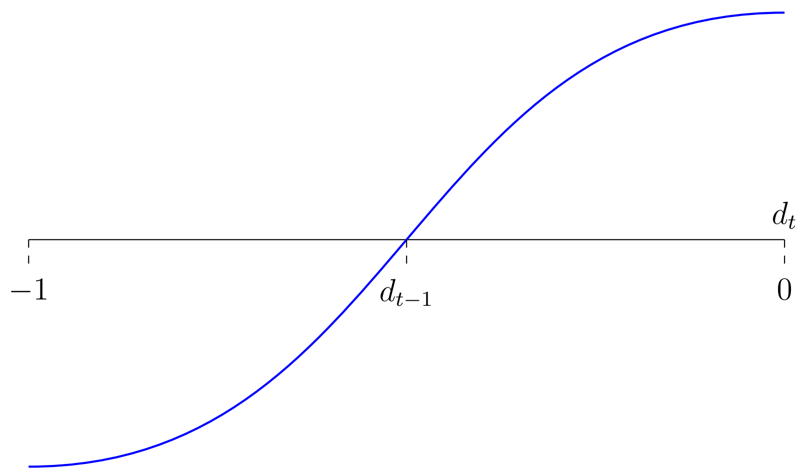

Figure 1: Anchor Function

policy to her left, $d_t < x$, we assume that the voter uses a (possibly degenerate) mixture of the two anchors: $u_x(d_t|d_{t-1}, r_{t-1}) = -|d_t - x| + \alpha f(d_t - d_{t-1}) + (1-\alpha) f(d_t - r_{t-1})$ for $\alpha \in [0, 1]$. Instead, for $r_t > x$ there is no anchor, as such we assume for simplicity that the anchoring function drops out of the voter's evaluation, $u_x(r_t) = -|r_t - x|$. We make the analogous assumptions when $x < d_{t-1} < r_{t-1}$. Note that when we turn to party competition, equilibrium platforms are on either side of the median voter, and her voting decision is unaffected by this assumption (which can be replaced by any assumption that preserves Lemma 1, and thus median decisiveness).

Remark. Formalizing the use of policy anchors requires us to make assumptions on how they are used. In our theory, voter anchoring is direction-specific. Voters compare today's platform to her left to yesterday's platform that was to her left, and not to a polity-wide variable such as the last election's outcome, or an individual-specific variable such as a person's current wealth level. As discussed in the introduction, this approach is consistent with the direction-specific contextual evaluations found by Simonovits (2017). For simplicity, anchoring is only to the platform from the previous election rather than the full policy history. Although one can imagine the effect of long-past platforms lingering in voters' memories, recency bias points toward the most recent platform being the dominant anchor point. In Section 5 we discuss the implications of allowing the voters to have a longer memory.

2.2 Implications for Voting Behavior

Policy anchors do not upset the basic principle of spatial voting that, all else equal, a policy platform closer to the voter is more favorable than a platform further away. Furthermore, they do not overturn the standard single-crossing property, i.e., policy preferences are well-ordered across voters. Lemma 1 formalizes this property.

Lemma 1 *$u_x(p_t|p_{t-1})$ is single-peaked in p_t, with unique maximizer at the voter's ideal policy, x. Moreover, it satisfies the strict single-crossing property in (x, p_t).*

Policy anchors *do*, however, violate the standard implication of spatial voting that a voter votes for the policy platform that is closest to her ideal. To see this violation, consider the situation depicted in Figure 2 of a voter deciding between two parties, D and R. Party R's platform is closer to the voter's ideal policy at x than is party D's. However, party R has not moved from its platform in the previous election whereas party D has moved toward the voter.

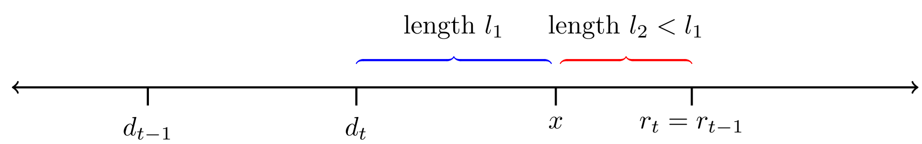

Figure 2: Within-Election Preference Reversal

Party D's shift has made it more appealing to the voter. If the boost from the shift is significant enough, it can induce the voter to support D over R despite the greater distance of D's platform in this election. Such a switch represents a preference reversal relative to the standard spatial voting theory.[9] We refer to it as the *within-election* preference reversal.

Within-Election Preference Reversal A voter may vote for a party other than the one with a platform closest to her ideal policy.

This same effect can lead to a preference reversal across elections. Even if the parties do not change their platforms from one election to the next, a voter may switch her vote from one party to the other. This possibility is depicted in Figure 3 that continues the example from Figure 2 to one more election. From election $t-1$ to election t, D's initially extreme position gave it a boost in the eyes of the voter at election t. But after election t, the anchors reset, and the boost disappears, such that D loses this individual's vote at $t+1$ because it is further from the voter than party R. We refer to this as the *across-elections* preference reversal.

[9] A voter with asymmetric preferences around her ideal policy may also prefer d_t over r_t, although the spatial violation would go away with an appropriate rescaling of the policy space. Preference reversals due to the use of anchors are possible with any scaling of the policy space.

Across-Elections Preference Reversal For pairs of policy platforms that are unchanged from one election to the next, $d_{t+1} = d_t$ and $r_{t+1} = r_t$, a voter may switch which party she votes for from election t to election $t+1$.

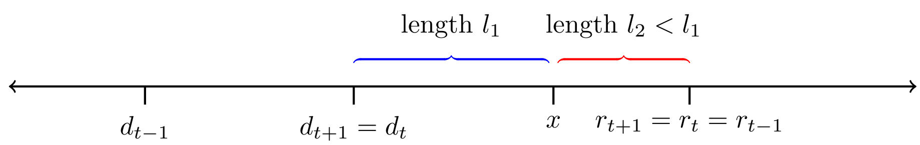

Figure 3: Across-Elections Preference Reversal

3 Electoral Competition with Anchors

We turn now to how voters' use of anchors impacts politics more broadly. It is natural that those competing for voters' affection consider changes in voters' preferences, and thus adjust and adapt to the use of anchors. To study these follow-on effects, we develop a model of electoral competition and characterize how voters' use of anchors affects the policy platforms offered by political parties.

The model builds on the classic Calvert-Wittman framework with policy-motivated parties, extending it dynamically. To highlight the role of anchors, we otherwise remain faithful to their framework.

In each period $t = 1, 2, ...$, two parties D and R compete in an election by simultaneously offering platforms, d_t and r_t, respectively, in the policy space $[-1, 1]$. The anchor for a party is its previous policy offering, $d_{t-1} \leq 0$ for party D and $r_{t-1} \geq 0$ for party R. The anchors for the first election, d_0 and r_0, are exogeneous, whereas they are endogenous for all later elections. Voters observe the proposals, cast their votes in a majority rule election, and the winning party implements its platform.

By Lemma 1, voter utility functions are single-peaked over $[-1, 1]$ and satisfy the single-crossing property in policies and ideal points. Therefore, the median voter is decisive over pairwise comparisons of policy platforms (Rothstein, 1991; Gans and Smart, 1996). As such, moving forward we focus on the vote of the median voter (or "the voter"), and we assume she has an ideal policy in the center of the policy space at 0.

In addition to policy, the voter cares about the valence of each party.[10] Normalizing the valence of the left party to 0, R's valence is given by $\theta_t \sim U(-\psi, \psi)$. The valence θ_t is

[10]The valence shock is common and does not affect Lemma 1.

unknown to the parties until after they have chosen their policy platforms, but it is realized before voting occurs.

Thus, the voter's utility for electing D is

$$u_0(d_t|d_{t-1}) = -|d_t| + f(d_t - d_{t-1}), \tag{3}$$

while the voter's utility from electing R is

$$u_0(r_t|r_{t-1}) + \theta_t = -|r_t| + f(r_{t-1} - r_t) + \theta_t. \tag{4}$$

The parties are motivated purely by policy outcomes. The ideal policy of party D is -1 and of party R is $+1$. Their utility functions from a policy p are

$$v_D(p) = -|p+1| \tag{5}$$

$$v_R(p) = -|p-1|. \tag{6}$$

Intertemporal linkages. The key linkage between elections is that the anchor points evolve endogenously: today's platforms becomes tomorrow's anchor points. To focus on this, we shut down other linkages across time. Specifically, we assume that the parties (and voters) are myopic, so that they do not consider how today's platform choice will influence their utility in future periods. We revisit and discuss the role of foresight later in Section 4.2. Similarly, the valence draw in each election is independent, such that there is no learning about voter preferences over time.

The electoral environment. In practice, the electoral environment evolves over time as new issues emerge, the salience of issues changes, and voter demographics and preferences update with changing times and circumstances. To understand the full impact of voter anchoring, we thus develop a model that allows for change to the environment. For focus and tractability, we limit this change to the valence term. Specifically, we index ψ by the time period t and we suppose that $\psi_t \in \{\psi_l, \psi_h\}$, with $\psi_l < \psi_h$, evolves as a Markov chain with a transition probability $\rho \in (0,1)$.[11] This means that the state is either ψ_l or ψ_h and in each period it switches to the other state with probability ρ. With probability $1-\rho$ the state is unchanged. A small ρ, therefore, represents a very stable environment. We assume that ψ_t is drawn and publicly observed at the start of each subsequent period.

To build intuition we begin with the case where the environment is unchanging, i.e., ψ_t is fixed at the same ψ in every period. As such, we will drop the time subscript from ψ when there is no room for confusion.

[11] In Section 5 we discuss how our results extend if ψ_t is drawn from a distribution with continuous support.

The anchoring function. Finally, we impose the following technical conditions on the anchoring function and support of the valence shock:

Assumption 1 *For any Δp and ψ we assume:*

1. $1 + 2f(1) < \psi$.
2. $\psi < 2$ *and* $f(1) < 1/2$.
3. $\frac{f''(\Delta p)}{1+f'(\Delta p)} < 1$.

Part 1 ensures that there is always an interior probability of winning the election. Part 2 implies that the standard Calvert-Wittman equilibrium is interior to the parties' ideal points, while still being consistent with part 1. Finally, part 3 guarantees existence and uniqueness of a symmetric equilibrium.

4 Results

We begin by describing the equilibrium party platforms *within* an election and then turn to platform dynamics *across* elections.

4.1 Party Platforms Within an Election

Consider the first election at $t = 1$. The voter votes for party R if $u_0(r_1|r_0) + \theta_1 \geq u_0(d_1|d_0)$, and for party D otherwise. Given her ideal point of 0, and that θ_1 is drawn uniformly over $[-\psi, \psi]$, the probability R wins the election given platforms (r_1, d_1) and anchors (r_0, d_0) is:

$$P(r_1, d_1|r_0, d_0) \equiv \frac{1}{2} - \frac{r_1 + d_1 + f(d_1 - d_0) - f(r_0 - r_1)}{2\psi}.$$

As the parties are policy motivated, they balance the probability of winning against the policy outcome. Party R's maximization problem is described by:

$$\max_{r_1} P(r_1, d_1|\cdot)v_R(r_1) + \big(1 - P(r_1, d_1|\cdot)\big)v_R(d_1). \tag{7}$$

Party D's problem is analogous. As is standard, we focus on symmetric equilibria. In our setting, this requires the additional condition that the initial (exogenous) anchors are symmetric as well, $r_0 = -d_0$. Given this, a unique symmetric equilibrium exists and is given by the following.

Proposition 1 *There exists a unique symmetric equilibrium* $r_1^* = -d_1^* \in (0, 1)$ *where*

$$r_1^* = \frac{\psi}{2\Big(1 + f'(r_0 - r_1^*)\Big)}. \tag{8}$$

The effect of the policy anchor is evident in the denominator of Equation (8). This result encompasses the canonical model in which the anchor term disappears (equivalent to the slope of $f(\cdot)$ being zero everywhere). In that case, the equilibrium in Proposition 1 reduces to the result of Calvert and Wittman.

Corollary 1 *Without voter anchoring, the equilibrium platforms are* $r_{cw}^* = -d_{cw}^* \in (0, 1)$ *where*

$$r_{cw}^* = \frac{\psi}{2}. \tag{9}$$

The equilibrium without anchors is independent of the history. Thus, wherever the parties were previously located, they jump to the locations given by Equation (9).

When voters use policy anchors, equilibrium locations depend on the history. A first observation is that anchoring provides a convergent force on party locations. To see this, recall the classic incentives of policy motivated parties: By moderating toward the median voter, a party increases its probability of winning but wins with a less attractive policy. Without policy anchoring, these forces are in perfect balance for party R at r_{cw}^*. With policy anchoring, and regardless of where the anchor is, moderating towards the median gives R a slightly larger boost in its probability of winning as the new platform will look more attractive (or less unattractive) relative to the anchor. This is evident formally in Equation (8) as the slope of $f(\cdot)$ is everywhere strictly positive. It follows that the equilibrium with anchoring is, all else equal, always more convergent than it is without anchoring.

Corollary 2 *The equilibrium platforms with voter anchoring satisfy* $r_1^* \in (0, r_{cw}^*)$.

Despite this result, it does not follow that the parties always moderate in order to look better to voters than they did previously. For anchors that are very moderate, the parties will in fact adopt platforms that are more extreme than their anchors, and they do this despite knowing that the comparison will be viewed poorly by voters.

To see why, return again to the classic trade-off parties face, this time turning it around: by choosing a more extreme platform, a party lowers its probability of winning but wins with a more appealing policy. This incentive is strong when starting from a very moderate position, so even though a more extreme platform comes with the additional cost of a poor comparison to the policy anchor, it can still be worth it.

A key threshold in the logic of policy positioning in this setting is the anchor that is the fixed point of Equation (8), which we denote by $\bar{r}$. When the policy anchor for party R is at $\bar{r}$, its optimal choice is $\bar{r}$ itself. Formally, $\bar{r}$ satisfies $r_1^*(r_0) = r_0$, and is explicitly given by:

$$\bar{r} = \frac{\psi}{2\Big(1 + f'(0)\Big)}. \tag{10}$$

Recall that the anchoring function is s-shaped, i.e., f' is maximized at 0. As a consequence, $\bar{r}$ is a lower bound on the the equilibrium platforms of the right-wing party. Proposition 2 describes how the equilibrium varies in the location of the anchor around $\bar{r}$.

Proposition 2 *If:*

1. *$r_0 < \bar{r}$ then $r_1^* > \bar{r}$, and decreasing in r_0.*
2. *$r_0 > \bar{r}$ then $r_1^* \in (\bar{r}, r_0)$, and is increasing in r_0 at rate less than one.*
3. *$r_0 = \bar{r}$ then $r_1^* = \bar{r}$.*

Figure 4 depicts the equilibrium. For an anchor more moderate than $\bar{r}$, the equilibrium platforms are more extreme than the anchor, whereas for anchors more extreme than $\bar{r}$ the platforms are more moderate. Furthermore, the equilibrium platforms relative to the anchor are not symmetric around $\bar{r}$. For moderate anchors, the platforms are substantially more extreme than the anchor, with the gap decreasing the closer the anchor is to $\bar{r}$. For more extreme anchors, instead, the platforms are incrementally more moderate than the anchor, with the gap increasing the more extreme is the anchor. As such, the equilibrium policy platforms are non-monotonic in the anchor, with the most extreme policy positions coming from moderate and extreme anchors, and the least extreme policy platforms arising when the anchor is $\bar{r}$, and Party R locates exactly there. In other words, the equilibrium platforms respond to the anchors in a U-shaped manner.

This pattern is the consequence of the decreasing marginal effect of moving away from the anchor. Small differences with the anchor lead to relatively large marginal benefits and costs whereas larger differences lead to relatively smaller marginal benefits and costs. Thus, when the anchor is extreme and the parties wish to offer more moderate policies, their incentives are tempered by a marginal anchoring boost that is decreasing in the relative movement, and so the parties only make small incremental shifts towards the voter. Symmetrically, the marginal cost of shifting away from the median voter is decreasing in the size of the shift. Starting from centrist anchors, the parties shift outward more than they otherwise would, overshooting the threshold $\bar{r}$. The asymmetric reaction of platforms to the anchor reflects this logic, forming

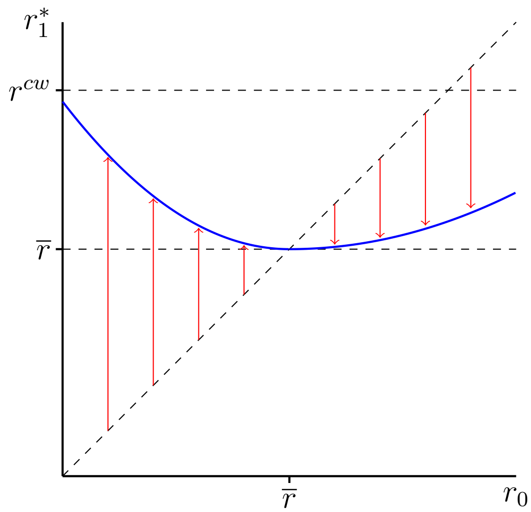

Figure 4: Equilibrium Platforms

a single-dipped function such that centrist and extreme anchors give rise to local maxima in the equilibrium platforms. The overshooting of $\overline{r}$ when starting from a moderate anchor is a key insight of our model. It is the driving force of cycles in party platforms, to which we now turn.

4.2 The Dynamics of Party Platforms

Over time, the policy anchors evolve with political competition. Voters use the previous positions of the parties to evaluate their current platforms, and the current platforms then become the anchors for the following election. In this section we study the dynamics in party platforms that this recursive process generates. We begin with a political environment that is unchanging (ψ_t is fixed) and then turn to the more realistic setting in which the environment itself changes over time (ψ_t may fluctuate between ψ_l and ψ_h).

4.2.1 Party Platforms in an *Unchanging* World

In an unchanging environment, the competitive incentives of the parties recur from one election to the next. Without policy anchors, equilibrium behavior is straightforward: In the first election the parties locate at d_{cw} and r_{cw} and remain there for each election thereafter.

With policy anchors, equilibrium behavior is very different. Even in an unchanging world, the policy platforms trace a rich dynamic of positions, changing from one election to the next, possibly reversing direction, and converging toward but never settling down to stable locations. It is only if the initial anchor is precisely at the threshold $\overline{r}$ that equilibrium

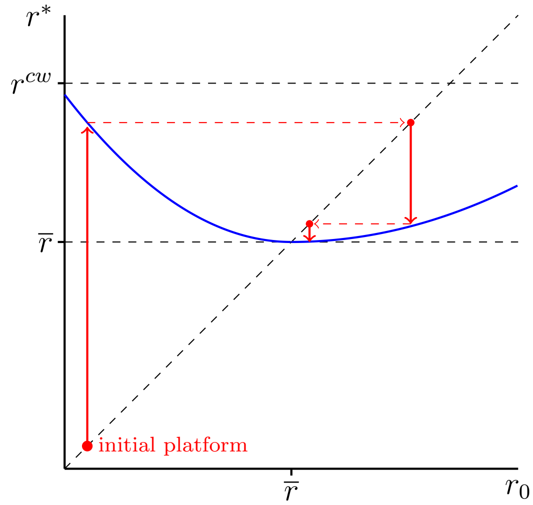

Figure 5: Equilibrium Platform Dynamics

platforms are ever stable.

Proposition 3 characterizes policy dynamics in this setting.

Proposition 3 *In the election at $t = 1, 2, ...,$ if:*

1. *$r_0 > \overline{r}$ then $r_t^* \in (\overline{r}, r_{t-1})$ for all t, and $r_t^* \to \overline{r}$ as $t \to \infty$.*

2. *$r_0 < \overline{r}$ then $r_1^* > \overline{r}$, and dynamics then follow case (1).*

3. *$r_0 = \overline{r}$ then $r_t^* = \overline{r}$ for all t.*

Figure 5 traces out a sample path of the platform dynamics generated by the behavior in Proposition 3. Starting from a moderate anchor, the red arrows and dashed lines trace through the equilibrium path of R's policy platform. Initially, R's platform makes a large jump outward as the shape of the anchoring function means that the marginal cost of moving away from the voter is decreasing. From there, R's platform creeps inward, with R benefiting each election from a platform that is appealing relative to the anchor. R's platform converges toward the threshold $\overline{r}$ but its steps get smaller at a faster rate and its platform never quite reaches $\overline{r}$. Thus, equilibrium platforms initially jump outward, across the $\overline{r}$ threshold, before reversing direction and iterating toward $\overline{r}$ but never reaching it.

Reversals in the direction of party platforms represent a time-inconsistency in party incentives that is due to voter anchoring. The parties have an incentive to diverge starting from a moderate position to try and win with a better policy, i.e., the usual Calvert-Wittman logic. Recall that the marginal cost of moving outward is decreasing due to the S-shaped anchoring function. Thus, the parties jumps out from a moderate position, overshooting $\overline{r}$. Once at the new location, however, the anchor resets and the low marginal cost of divergence disappears. Instead there is a high marginal anchoring benefit for small steps inward. Consequently, the party shifts inward to a moderate policy, but does so in small steps to avoid the diminishing marginal benefits of a large jump. It is only at $\overline{r}$ that party R's policy incentive to diverge to a better policy is exactly in balance with the anchoring benefit of converging.

Two remarks on the foresight of parties. First, it is worth emphasizing that the behavior in Proposition 3 emerges even though parties are *not* forward-looking. When a party moves away from the median voter in equilibrium, it is not that it is suffering a cost today to benefit from a more attractive anchor in the future. In each election the parties myopically optimize, thus a party shifts outward because there is an immediate policy benefit today that is worth the lower probability of winning that the shift causes. The shifts towards and away from the median voter, as well as their asymmetry, derive from voter anchoring and not forward-looking strategic behavior from the parties.

Second, although beyond the scope of this paper, how equilibrium platforms are affected by forward-looking parties is a question of natural interest. A reasonable conjecture is that a qualitatively similar dynamic would emerge. Indeed, intuition suggests that the nature of platform shifts inward and outward would be exaggerated by forward-looking behavior because the parties would seek to exploit dynamically the asymmetry in how voters respond to anchors. The parties would benefit from smoothing across time the large marginal benefits of creeping inward (Figure 5), thus slowing down convergence of their platform to the stable position, and may overdo the jumps outward to leave even more space to accrue these marginal benefits.[12]

4.2.2 Party Platforms in a *Changing* World

In this section we allow for the environment itself to change over time, where the state of the world is the support of the valence shock, either ψ_l or ψ_h. The size of ψ_t affects the incentives of parties to converge or diverge. For the smaller value, ψ_l, an incremental movement toward the center increases the probability of winning by more than for the high value, ψ_h. This incentive, in turn, leads to a lower value of the threshold $\overline{r}$, and we denote the level associated with each value of $\psi_t \in \{\psi_l, \psi_h\}$ by $\overline{r}_l$ and $\overline{r}_h$, where $\overline{r}_l < \overline{r}_h$.

Without policy anchors, equilibrium behavior is once again straightforward. For each state, the equilibrium platforms are constant, and when the state changes, the equilibrium shifts instantly and completely from one value to the other. When the state switches from ψ_l to ψ_h the platforms shift from moderate to extreme, and vice versa when the state switches

[12]See Izzo (2023) for a model where a similar dynamic incentive emerges.

back. The platforms shift, therefore, only when the state changes, and by a fixed and constant amount each time. This is depicted in Figure 6b.

With policy anchors, the platform dynamics are much richer. The platforms never stabilize. They endlessly evolve following an asymmetric pattern, with incremental shifts and moderate jumps towards the median voter, interspersed with larger occasional jumps away from her. This pattern holds even if the state changes infrequently and there is little difference in the states, i.e., $\rho > 0$ and $\psi_h - \psi_l > 0$ can be arbitrarily small. In an environment that is mostly stable, voter anchoring causes platforms to change in a constant flow, cycling between moderate and extreme positions.

Proposition 4 describes the equilibrium behavior distinguishing between platform change when the state changes that period and when it remains the same.

Proposition 4 *Suppose $r_0 > \overline{r}_h$. In the election at $t = 1, 2, ...$, the following holds:*

- *If the state does not change at time t, $\psi_t = \psi_{t-1}$, then $r_t^* < r_{t-1}^*$, and $|r_{t-1}^* - r_t^*|$ is increasing in r_{t-1}^*.*
- *If the state changes at time t from $\psi_{t-1} = \psi_h$ to $\psi_t = \psi_l$, then $r_t^* \in (\overline{r}_l, r_{t-1}^*)$ and $|r_{t-1}^* - r_t^*|$ is increasing in r_{t-1}^*.*
- *If the state changes at time t from $\psi_{t-1} = \psi_l$ to $\psi_t = \psi_h$ and $r_{t-1}^* < \overline{r}_h$, then $r_t^* > \overline{r}_h > r_{t-1}^*$, and $|r_{t-1}^* - r_t^*|$ is decreasing in r_{t-1}^*. If $r_{t-1}^* > \overline{r}_h$, then $r_{t-1}^* > r_t^* > \overline{r}_h$ and $|r_{t-1}^* - r_t^*|$ is increasing in r_{t-1}^*.*

The path of platforms over time is depicted in Figure 6a. If a state persists from one election to the next, the party platforms creep inward toward the threshold $\overline{r}_l$ or $\overline{r}_h$ with convergence slowing over time. Therefore, the size of the movement is increasing in the anchor. When the state shifts from ψ_h to ψ_l, the platform continues moving in the same direction although it makes a larger jump inward. Likewise, if the state switches from ψ_l to ψ_h and the anchor is greater than $\overline{r}_h$, which is only true if the state has not been at ψ_l for very long, then platforms also continue to move towards the center even after the state changes, albeit at a slower rate than before.

Balancing this inward movement is the response when the state shifts from ψ_l to ψ_h and the anchor is below $\overline{r}_h$, which must eventually happen if platforms have only been moving inward. When $\psi = \psi_h$ the equilibrium platform is always above the cutoff $\overline{r}_h$. Thus, after the state changes, platforms reverse course, making a large jump away from the center, as shown in Figure 6a. With the anchor now more extreme than $\overline{r}_h$, the process renews itself and convergence begins anew. This process thus generates stochastic, but predictable, cycles

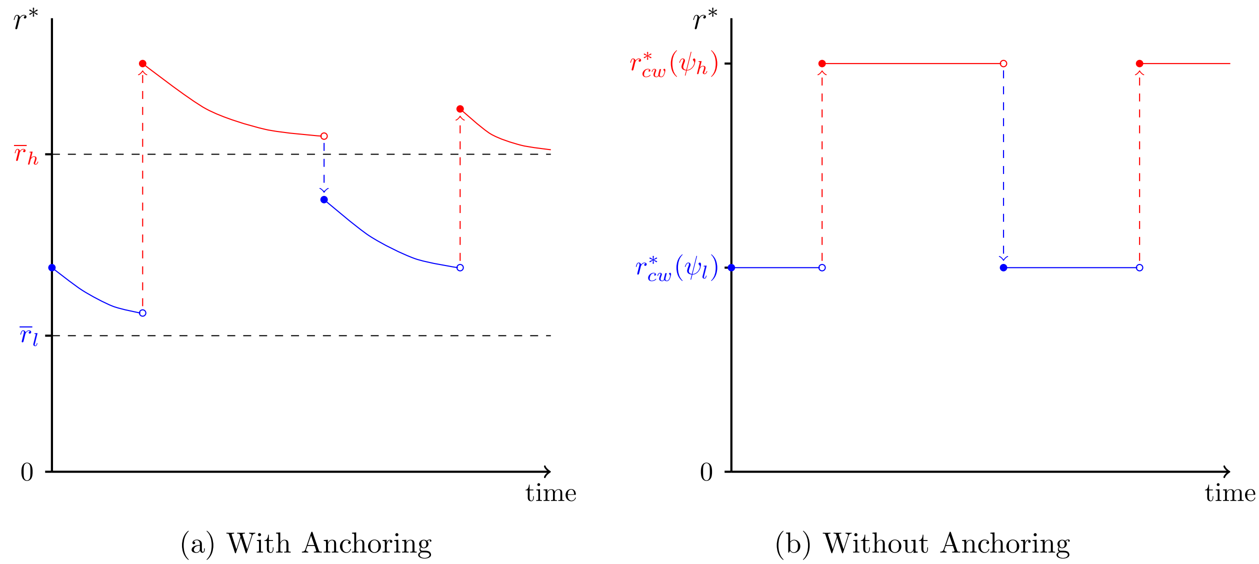

Figure 6: Platform Dynamics in a Changing World

in party platforms. In sharp contrast with the world without anchoring (see Figure 6b), the size of the platform jumps and the amplitude of the dynamic path is not a constant, rather it varies with the position of the anchor.

Notice that an implication of this result is that the magnitude of platform movement varies in the number of elections between a change in state, and it depends on whether platforms continue moving in the same direction or not. In particular, consider periods in which the inherited state is $\psi_{t-1} = \psi_l$, as these are the periods in which we may see a reversal in the direction of platform movement. The more time the political environment spends in state ψ_l the further R's platform moderates towards 0 and away from $\overline{r}_h$. Consequently, when the state eventually switches to ψ_h, the anchor is far from $\overline{r}_h$ and the platform makes a large jump outward. In this way, the asymmetric changes around the $\overline{r}$ threshold in an unchanging world manifest in a changing world as asymmetric and history dependent sized jumps when the environment switches from one state to the other.

Corollary 3 *Assume $\psi_t = \psi_l$ for all periods $t \in \{1, ..., T-1\}$ and $r^*_{T-1} < \overline{r}_h$.*

- *If $\psi_T = \psi_L$ we have $r^*_T < r^*_{T-1}$. Then $|r^*_{T-1} - r^*_T|$ is decreasing in T.*
- *If $\psi_T = \psi_H$ we have $r^*_T > r^*_{T-1}$. Then $|r_{T-1} - r^*_T|$ is increasing in T.*

This result highlights that stability is elusive in our setting. If the state remains at ψ_l over a long period time, then platforms will appear to be stabilizing by making smaller and smaller movements towards the center. However, it is precisely under these conditions that the platforms change most dramatically when the state eventually shifts and they reverse course.

While the previous two results study how the location of the anchor determines the path of policy platforms, our last result builds on these dynamics to draw empirical implications about the nature of changes in parties' policy offerings. Specifically, as the parties continue to compete over a long period of time, their equilibrium behavior will generate a stationary distribution of observed platforms. Proposition 5 provides predictions about changes in extremism based on this distribution of platforms.

Proposition 5

- *The probability that we observe platforms become more moderate is greater than the probability they become more extreme.*
- *For δ sufficiently large, the probability that we observe platforms become more extreme by at least δ is greater than the probability we observe platforms become more moderate by at least δ.*

The model predicts that we should more often see platforms become more moderate, rather than more extreme. However, if we do see a large change in party platforms, it is more likely to be due to the parties moving to the extremes. Note that both of these predictions differ from the Calvert-Wittman model. Under the Calvert-Wittman model, we should observe half of the platforms at $r^*_{cw}(\psi_h)$ and half of the platforms at $r^*_{cw}(\psi_l)$. As such, it predicts that the probability platform extremism increases is equal to the probability that it decreases, and that the magnitude of changes is the same, regardless of the direction of movement in party platforms.

5 Concluding Discussion

In concluding, we discuss the implications of our theory and questions that it gives rise to. But first we address informally several variations and extensions of the model.

Variations and Extensions

More policy anchors. In the model, we assume that voters look only to the most recent platform of each party to evaluate the current proposals. One may instead imagine that a longer history of the platforms influences the voters' evaluation. In the context of our model, this can be captured, for example, by assuming that voters anchor their evaluation of today's proposals to a weighted average of past platforms over a finite history, with the weights potentially decreasing as older platforms fade in influence. In this case, voters' longer memory creates a lag in how the anchoring function responds to the platforms proposed along the path. This lag slows down the platform dynamics emerging in our baseline model, however, it does not need to alter their qualitative features. In particular, the same dynamics emerge in the long run, but they may be preceded by an initial period of seemingly more erratic movement, during which platforms shift back and forth between more extreme and more moderate positions before settling into the patterns we characterized above.

Another possibility is that voters only remember the most extreme (or most moderate) platforms proposed by the parties over some finite past number of periods. As in the baseline, the parties' platforms evolve with long periods of small movements towards the voter punctuated by larger jumps in the opposite direction. In contrast to the baseline model, however, temporary periods of stability emerge in this version of the game. This stability occurs because the same past policy may remain the most extreme platform for several periods in a row, until voters eventually use the more recent platform as an anchor. Interestingly, stability has an asymmetric feature, as only moves in one direction are followed by a finite number of periods of no change, after which the platforms resume their path as in the baseline model. Thus, similar to the previous case, voters forming anchors based on a longer history of platforms can slow the dynamics we characterized earlier, but does not overturn the long-run evolution of policies.

Asymmetric equilibria. For tractability, our analysis fixed the initial policy anchors for both parties symmetrically around the median voter's bliss point. Doing so allows us to focus on symmetric equilibria, and makes our findings comparable to earlier work on electoral competition which often focuses on symmetric equilibria as well. In our setting, however, a natural question is how platform dynamics would change if parties started from asymmetric anchors. Here, parties face similar incentives as in the baseline model: to moderate when their anchor is extreme, and move more extreme when their anchor is moderate. As such, platform dynamics in an asymmetric equilibrium should have many of the same features we discuss above. Platforms continuously evolve, and parties make smaller moves towards the median voter followed by larger jumps in the opposite direction.

However, two crucial differences from behavior in our symmetric equilibrium could emerge. First, in the symmetric equilibrium, it is always the case that *both* parties move more moderate or *both* parties move more extreme. In contrast, with asymmetric platforms and anchors there could be periods in which one party 'chases' the other. That is, a party starting from an extreme anchor moves toward the median voter, while its opponent, with a very moderate anchor, shifts away. Importantly, this dynamic would lead to a within-election preference reversal for the median voter on the equilibrium path — an outcome that cannot arise when parties are always positioned symmetrically. Second, even when the parties both become more extreme or both become more moderate, the extent of this movement could differ across parties. Thus, for example, we could observe both parties becoming more extreme, but with one party shifting significantly more than the other.

More than two states. To clearly illustrate our results, our model with a changing world only considers a binary state. If instead ψ_t is drawn from a continuous support in periods where the state changes, then in almost all states there is some probability of observing platform movement reverse direction, but the underlying logic and results remains unchanged. That is, for a given state ψ there is a cut-off $\overline{r}_\psi$ that determines whether party platforms moderate or polarize. Absent a state change, when the starting policy anchor is below $\overline{r}_\psi$ the party platforms jump outward and then begin to slowly creep inward toward $\overline{r}_\psi$, as in the baseline model. As such, when platforms initially begin moving more moderate, the anchors are very extreme. Thus, if there is a state change, $\psi_t \neq \psi_{t-1}$, it is very likely that the new $\overline{r}_{\psi_t}$ is still more moderate than the anchor, and platforms continue to creep inward. However, if platforms have repeatedly moved inward over a long period of time, then it becomes increasingly likely that a new draw of ψ_t yields a $\overline{r}_{\psi_t}$ that is above the policy anchor. When this happens, parties adopt platforms more extreme than $\overline{r}_{\psi_t}$, halting the inward movement, triggering a large inward shift away from the voter, and restarting the cycle.

Thus, just as in our two-state model, small shocks to the environment will drive platforms to cycle indefinitely in an asymmetric fashion. Moreover, the longer the state remains stable, the more dramatic the platform reversals will be in equilibrium.

The Dynamics of Party Platforms

The movements of political parties have long fascinated political scientists. In contrast with implications from classic models of electoral competition, empirical scholars often point to rich dynamics in the trajectory of electoral platforms, with parties moving in a seemingly chaotic way shifting back and forth between more moderate and more extreme positions (see, e.g., Budge et al., 2010). Understanding these dynamics — why and when parties make bold jumps versus incremental change, why they sometimes reverse course and other times continue moving in the same direction — is important to understanding the nature of political competition and policy outcomes.

Our model shows that a relatively small departure from classic models of spatial competition, that accounts for voters' difficulties in evaluating policies, provides a rich set for implications for understanding platform dynamics. The observed frequent movement of platforms and the whipsaw pattern it features can be accounted for by voters' use of policy anchors. Additionally, the use of anchors predicts that these patterns should be asymmetric, with party platforms making frequent small shifts in one direction punctuated by occasional dramatic changes in the other, consistent with the dynamics that are also observed in public policy (Baumgartner and Jones, 1993).

Moreover, our model delivers unique predictions about the exact features of these asymmetric dynamics: moves toward the center should be frequent and small, while shifts to the extremes are rarer and larger. Empirically investigating this asymmetry is a clear direction for future work.

Furthermore, our theory can help us understand the apparent mismatch between the volatility of party platforms and the relative stability of voters' policy opinions (Page and Shapiro, 1992; Druckman and Leeper, 2012). In our model, while voters' constructed preferences over the policy space evolve endogenously, the ideological location of their ideal points remains fixed. Our results thus resonate with observed patterns insofar as surveys of public opinion measure the latter, and not the former.[13] For instance, survey questions that ask respondents to place their political views on a left-right scale or assess whether a given level of spending or government intervention is excessive or insufficient aim to capture the location of their ideal points. Prompts asking voters to express the extent to which they approve or disapprove of different parties, instead, capture their constructed preferences.

By making this distinction explicit, our theory underscores the importance of considering how public opinion is measured when assessing both the extent and the consequences of voter preference change. Ignoring this distinction risks conflating stable ideological bliss points with evolving, context-dependent evaluations—leading to misguided interpretations of electoral dynamics.

Political Representation

In our model, voters' ideal policies—their *fundamental* preferences—are unchanging, yet the policy platforms offered by the parties do change. Moreover, they change in response to how voters evaluate policies—their *constructed* preferences. If political representation requires that parties represent the preferences of voters, is this or is this not representation?

The scope of this question extends well beyond our model. Indeed, it lies at the heart of the classic debate over whether true representation requires elected officials to act as delegates or as trustees (Pitkin, 1967). In the usual framing of the problem, it is differences in expertise

[13] While our model holds voters' ideal points fixed, if these primitives fluctuated over time, platform dynamics would reflect *and amplify* such changes, consistent with the empirical literature. In that scenario, platforms would evolve in response to changes in ideal points, as in the standard CW model, *and* in response to shifts in the relevant anchors, as in our model.

between voters and politicians that create a tension between whether representatives should act as delegates or trustees. Instead, in our model it is cognitive limitations on the part of voters that create a wedge between voters' behavior and their long-term interests. As such, it remains an open question whether policymakers are representing voter preferences even though the parties commit to platforms and have no private information.

The wedge in voters' interests in our framework turns on the difference between voters constructed preference—the decision utility they experience when voting—and voters' true or fundamental preferences that come from their ideal point. Which quantity should we employ to evaluate the extent of successful representation?

One argument is that the fundamental preference should be used as it more accurately measures the real and persistent experience that comes from policies in action, whereas decision utility is ephemeral and exists only in the electoral moment. The counter argument is that decision utility is more real to voters and measures the choice they make rather than something that happens to them indirectly. Moreover, voters' experience in the electoral moment determines whether in their minds the electoral outcome is correct and, thus, is the foundation for democratic legitimacy

In our model, the parties are self-interested and respond to voters' decision utility. As we saw, this leads to continuous cycles in party platforms where the platforms would otherwise be largely stable. The fortuitous irony, however, is that appealing to voters' constructed preferences induces parties to converge closer to the ideal point of the median voter—recall Corollary 2—and this increases voter welfare as measured by their fundamental preferences.

References

Abdellaoui, M. (2000). Parameter-free elicitation of utility and probability weighting functions. *Management science 46*(11), 1497–1512.

Alesina, A. and F. Passarelli (2019). Loss aversion in politics. *American Journal of Political Science 63*(4), 936–947.

Baumgartner, F. R. and B. D. Jones (1993). *Agendas and instability in American politics.* University of Chicago Press.

Boxell, L., M. Gentzkow, and J. M. Shapiro (2024). Cross-country trends in affective polarization. *The Review of Economics and Statistics 106*(2), 557–565.

Budge, I., L. Ezrow, and M. D. McDonald (2010). Ideology, party factionalism and policy change: An integrated dynamic theory. *British Journal of Political Science 40*(4), 781–804.

Callander, S. and J. C. Carbajal (2022). Cause and effect in political polarization: A dynamic analysis. *Journal of Political Economy 130*(4), 825–880.

Callander, S. and C. H. Wilson (2006). Context-dependent voting. *Quarterly Journal of Political Science 1*(3), 227–255.

Callander, S. and C. H. Wilson (2008). Context-dependent voting and political ambiguity. *Journal of Public Economics 92*(3-4), 565–581.

Calvert, R. L. (1985). Robustness of the multidimensional voting model: Candidate motivations, uncertainty, and convergence. *American Journal of Political Science*, 69–95.

Druckman, J. N. and T. J. Leeper (2012). Is public opinion stable? Resolving the micro/macro disconnect in studies of public opinion. *Daedalus 141*(4), 50–68.

Fishburn, P. C. (1982). Nontransitive measurable utility. *Journal of Mathematical Psychology 26*(1), 31–67.

Gans, J. S. and M. Smart (1996). Majority voting with single-crossing preferences. *Journal of public Economics 59*(2), 219–237.

Grillo, E. (2016). The hidden cost of raising voters' expectations: Reference dependence and politicians' credibility. *Journal of Economic Behavior & Organization 130*, 126–143.

Grillo, E. and C. Prato (2023). Reference points and democratic backsliding. *American Journal of Political Science 67*(1), 71–88.

Izzo, F. (2023). Ideology for the future. *American Political Science Review 117*(3), 1089–1104.

Kahneman, D. and A. Tversky (1979). Prospect theory: An analysis of decision under risk. *Econometrica 47*(2), 263–292.

Kőszegi, B. and M. Rabin (2006). A model of reference-dependent preferences. *The Quarterly Journal of Economics 121*(4), 1133–1165.

Levy, G. and R. Razin (2025). Political social learning: Short-term memory and cycles of polarization. *American Economic Review 115*(2), 635–659.

Lichtenstein, S. and P. Slovic (2006). *The Construction of Preference.* New York: Cambridge University Press.

Lockwood, B., M. Le, and J. Rockey (2024). Dynamic electoral competition with voter loss-aversion and imperfect recall. *Journal of Public Economics 232*, 105072.

Lockwood, B. and J. Rockey (2020). Negative voters? Electoral competition with loss-aversion. *The Economic Journal 130*(632), 2619–2648.

Page, B. I. and R. Y. Shapiro (1992). *The rational public: Fifty years of trends in Americans' policy preferences*. University of Chicago Press.

Page, S. E. (2006). Path dependence. *Quarterly Journal of Political Science 1*(1), 87–115.

Pierson, P. (2000). Increasing returns, path dependence, and the study of politics. *American Political Science Review 94*(2), 251–267.

Pitkin, H. F. (1967). *The Concept of Representation*. Berkeley: University of California Press.

Rothstein, P. (1991). Representative voter theorems. *Public Choice 72*(2), 193–212.

Rotter, G. S. and N. G. Rotter (1966). The influence of anchors in the choice of political candidates. *The Journal of Social Psychology 70*(2), 275–280.

Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics 69*(1), 99–118.

Simonovits, G. (2017). Centrist by comparison: Extremism and the expansion of the political spectrum. *Political Behavior 39*, 157–175.

Tversky, A. and D. Kahneman (1974). Judgment under uncertainty: Heuristics and biases: Biases in judgments reveal some heuristics of thinking under uncertainty. *science 185*(4157), 1124–1131.

Tversky, A. and D. Kahneman (2000). *Choices, values, and frames*. Cambridge University Press New York.

Tversky, A. and I. Simonson (1993). Context-dependent preferences. *Management Science 39*(10), 1179–1189.

Villani, C. (2008). *Optimal transport: old and new*, Volume 338. Springer.

Waismel-Manor, I. and G. Simonovits (2017). The interdependence of perceived ideological positions: Evidence from three survey experiments. *Public Opinion Quarterly 81*(3), 759–768.

Wang, A. H.-E. and F.-Y. Chen (2019). Extreme candidates as the beneficent spoiler? Range effect in the plurality voting system. *Political Research Quarterly 72*(2), 278–292.

Wittman, D. (1983). Candidate motivation: A synthesis of alternative theories. *American Political Science Review 77*(1), 142–157.

A Proofs

Proof of Lemma 1. To start, we show that $u_x(p_t|p_{t-1})$ is single-peaked at x. Consider any p_t, p'_t with $x < p_t < p'_t$. We show $u_x(p_t|r_{t-1}) > u_x(p'_t|r_{t-1})$. First, suppose that $d_{t-1} < x < r_{t-1}$. The result then follows by noting that: (1) $-|x-p_t| > -|x-p'_t|$; and (2) $f(r_{t-1}-p_t) > f(r_{t-1}-p'_t)$, since f is strictly increasing in Δp_t. Second, suppose $d_{t-1} < r_{t-1} < x$. Then the anchoring function drops out and the result follows immediately by $-|x-p_t| > -|x-p'_t|$. Finally, assume $x < d_{t-1} < r_{t-1}$. The result then follows by noting that: (1) $-|x-p_t| > -|x-p'_t|$; and (2) $\alpha f(d_{t-1}-p_t) + (1-\alpha) f(r_{t-1}-p_t) > \alpha f(d_{t-1}-p'_t) + (1-\alpha) f(r_{t-1}-p'_t)$ because f is strictly increasing in Δp_t. The analogous argument yields the result for $p'_t < p_t < x$.

Next, we prove that $u_x(p_t|p_{t-1})$ satisfies the single-crossing property. Specifically, suppose $p_t < p'_t$ and $u_x(p'_t|p_{t-1}) > u_x(p_t|p_{t-1})$. We show that, for all $x < x'$, $u_{x'}(p'_t|p_{t-1}) > u_{x'}(p_t|p_{t-1})$.

First, note that, by single-peakedness of u_x, we cannot have $x < p_t < p'_t$ and $u_x(p'_t|\cdot) > u_x(p_t|\cdot)$. Second, take $p_t < p'_t < x'$. By single-peakedness of $u_{x'}$ it follows immediately that $u_{x'}(p'_t|\cdot) > u_{x'}(p_t|\cdot)$. Thus, there is only one remaining case to consider: $p_t < x < x' < p'_t$. We break the analysis into several subcases depending on the locations of the anchors.

1. If $d_{t-1} < x < x' < r_{t-1}$ then $x < x' < p'_t$ yields:

$$
\begin{aligned}
u_{x'}(p'_t|r_{t-1}) &= -|p'_t - x'| + f(r_{t-1} - p'_t) \\
&> -|p'_t - x| + f(r_{t-1} - p'_t) = u_x(p'_t|r_{t-1}).
\end{aligned}
$$

   Additionally, $p_t < x < x'$ implies:

$$
\begin{aligned}
u_{x'}(p_t|d_{t-1}) &= -|p_t - x'| + f(p_t - d_{t-1}) \\
&< -|p_t - x| + f(p_t - d_{t-1}) = u_x(p_t|d_{t-1}).
\end{aligned}
$$

   Thus, if $u_x(p'_t|r_{t-1}) > u_x(p_t|d_{t-1})$ then $u_{x'}(p'_t|r_{t-1}) > u_{x'}(p_t|d_{t-1})$, as required.

2. If $x < d_{t-1} < x' < r_{t-1}$ then $x < x' < p'_t$ yields:

$$
\begin{aligned}
u_{x'}(p'_t|r_{t-1}) &= -|p'_t - x'| + f(r_{t-1} - p'_t) \\
&> -|p'_t - x| + \alpha f(d_{t-1} - p'_t) + (1-\alpha) f(r_{t-1} - p'_t) = u_x(p'_t|d_{t-1}, r_{t-1}),
\end{aligned}
$$

   because $d_{t-1} - p'_t < r_{t-1} - p'_t$ implies $f(d_{t-1} - p'_t) < f(r_{t-1} - p'_t)$. Additionally, $p_t < x < x'$ implies:

$$
u_{x'}(p_t|d_{t-1}) = -|p_t - x'| + f(p_t - d_{t-1}) < -|p_t - x| = u_x(p_t),
$$

because $p_t - d_{t-1} < 0$ implies $f(p_t - d_{t-1}) < 0$. Thus, if $u_x(p'_t|d_{t-1}, r_{t-1}) > u_x(p_t)$ then $u_{x'}(p'_t|r_{t-1}) > u_{x'}(p_t|d_{t-1})$, as required.

3. If $d_{t-1} < x < r_{t-1} < x'$ then $x < x' < p'_t$ yields:

$$\begin{aligned} u_{x'}(p'_t) &= -|p'_t - x'| \\ &> -|p'_t - x| + f(r_{t-1} - p'_t) = u_x(p'_t|r_{t-1}), \end{aligned}$$

because $r_{t-1} < p'_t$ implies $f(r_{t-1} - p'_t) < 0$. Additionally, $p_t < x < x'$ implies:

$$\begin{aligned} u_{x'}(p_t|d_{t-1}, r_{t-1}) &= -|p_t - x'| + \alpha f(p_t - d_{t-1}) + (1-\alpha) f(p_t - r_{t-1}) \\ &< -|p_t - x| + f(p_t - d_{t-1}) = u_x(p_t|d_{t-1}), \end{aligned}$$

because $f(p_t - r_{t-1}) < f(p_t - d_{t-1})$ by $p_t < d_{t-1} < r_{t-1} < x$. Thus, if $u_x(p'_t|r_{t-1}) > u_x(p_t|d_{t-1})$ then $u_{x'}(p'_t) > u_{x'}(p_t|d_{t-1}, r_{t-1})$, as required.

4. If $d_{t-1} < r_{t-1} < x < x'$ then $x < x' < p'_t$ yields:

$$u_{x'}(p'_t) = -|p'_t - x'| > -|p'_t - x| = u_x(p'_t),$$

Additionally, $p_t < x < x'$ implies:

$$\begin{aligned} u_{x'}(p_t|d_{t-1}, r_{t-1}) &= -|p_t - x'| + \alpha f(p_t - d_{t-1}) + (1-\alpha) f(p_t - r_{t-1}) \\ &< -|p_t - x| + \alpha f(p_t - d_{t-1}) + (1-\alpha) f(p_t - r_{t-1}) = u_x(p_t|d_{t-1}, r_{t-1}), \end{aligned}$$

Thus, if $u_x(p'_t) > u_x(p_t|d_{t-1}, r_{t-1})$ then $u_{x'}(p'_t) > u_{x'}(p_t|d_{t-1}, r_{t-1})$, as required.

5. If $x < x' < d_{t-1} < r_{t-1}$ then the result follows from a similar argument as in Case 5.

6. If $x < d_{t-1} < r_{t-1} < x'$ then $x < x' < p'_t$ yields:

$$u_{x'}(p'_t) = -|p'_t - x'| > -|p'_t - x| + \alpha f(d_{t-1} - p'_t) + (1-\alpha) f(r_{t-1} - p'_t) = u_x(p'_t|d_{t-1}, r_{t-1}),$$

which follows because $f(d_{t-1} - p'_t) < f(r_{t-1} - p'_t) < 0$ by $x < d_{t-1} < r_{t-1} < p'_t$. Additionally, $p_t < x < x'$ implies:

$$\begin{aligned} u_{x'}(p_t|d_{t-1}, r_{t-1}) &= -|p_t - x'| + \alpha f(p_t - d_{t-1}) + (1-\alpha) f(p_t - r_{t-1}) \\ &< -|p_t - x| = u_x(p_t), \end{aligned}$$

which follows because $f(p_t - r_{t-1}) < f(p_t - d_{t-1}) < 0$ by $p_t < d_{t-1} < r_{t-1} < x'$. Thus, if $u_x(p'_t|d_{t-1}, r_{t-1}) > u_x(p_t)$ then $u_{x'}(p'_t) > u_{x'}(p_t|d_{t-1}, r_{t-1})$, as required.

■

Lemma A1 *R's probability of winning in each election is interior, $P(r_t, d_t|\cdot) \in (0, 1)$.*

Proof. First, we show $P(r_t, d_t|\cdot) < 1$. This requires

$$\frac{1}{2} - \frac{r_t + d_t + f(\Delta d_t) - f(\Delta r_t)}{2\psi} < 1$$
$$\Leftrightarrow 0 < r + d + f(\Delta d_t) - f(\Delta r_t) + \psi$$

Note that $r_t \geq 0$ and $d_t \geq -1$. Therefore, a sufficient condition for the above to hold is that:

$$0 < 0 - 1 + f(\Delta d_t) - f(\Delta r_t) + \psi$$
$$\Leftrightarrow 1 < f(\Delta d_t) - f(\Delta r_t) + \psi. \tag{11}$$

Finally, notice that $f(\Delta d_t) \geq f(-1 - 0)$ and $f(\Delta r_t) \leq f(1 - 0)$. Thus, inequality (11) holds if:

$$1 < f(-1 - 0) - f(1 - 0) + \psi$$
$$\Leftrightarrow 1 + 2f(1) < \psi,$$

where the second line follows from the assumption that $f(a) = -f(-a)$ and the last inequality holds by Assumption 1.

Second, we show that $P(r_t, d_t|\cdot) > 0$. This requires

$$\frac{1}{2} - \frac{r_t + d_t + f(\Delta d_t) - f(\Delta r_t)}{2\psi} > 0$$
$$\Leftrightarrow \psi > r_t + d_t + f(\Delta d_t) - f(\Delta r_t),$$

which again holds by assumption that $1 + 2f(1) < \psi$. ■

Lemma A2 *For any anchors (r_{t-1}, d_{t-1}), if (r_t^*, d_t^*) is an interior equilibrium then it solves:*

$$\Big(\frac{-1 - f'(r_{t-1} - r_t)}{2\psi}\Big)(r_t - d_t) + P(r_t, d_t|r_{t-1}, d_{t-1}) = 0 \tag{12}$$
$$\Big(\frac{1 + f'(d_t - d_{t-1})}{2\psi}\Big)(r_t - d_t) - 1 + P(r_t, d_t|r_{t-1}, d_{t-1}) = 0 \tag{13}$$

Proof. Follows by differentiating R's expected utility in (7) with respect to r_1, and likewise for the analogous problem for D. ■

Lemma A3 *Party i's expected utility is strictly concave in its own policy choice.*

Proof. We prove the result for party R, the case of party D follows symmetrically. Fixing party D's policy at d_t and differentiating yields:

$$\Big(\frac{-1-f'(r_{t-1}-r_t)}{2\psi}\Big)(r_t-d_t)+1-\frac{r_t+d_t+f(d_t-d_{t-1})-f(r_{t-1}-r_t)+\psi}{2\psi}. \tag{14}$$

Differentiating again and rearranging, the second derivative of R's expected utility is negative if:

$$f''(\Delta r_t)(r_t-d_t)<2(1+f'(\Delta r_t)). \tag{15}$$

In each period, $r_t, d_t \in [-1,1]$, therefore, $r_t-d_t \leq 2$. Thus, a sufficient condition for inequality (15) to hold is that $2f''(\Delta r_t)<2\big(1+f'(\Delta r_t)\big)$, which holds by Assumption 1. ■

Proof of Proposition 1. We first show that any symmetric equilibrium must be interior, $r_t^* \in (0,1)$. By the standard argument in the Calvert-Wittman setting we cannot have an equilibrium with $r_t^*=0$. Next, suppose that $r_t^*=1$ is an equilibrium. Fix D's policy at $d_t=-1$ and $d_{t-1}=-r_{t-1}$, and consider R's best response. A necessary condition for $r_t=1$ to be a best response is that R's utility is increasing as $r_t \to 1$. That is,

$$\lim_{r_t\to 1}\Big(\frac{-1-f'(\Delta r_t)}{2\psi}\Big)(r_t+1)+1-\frac{r_t-1+f(-1+r_{t-1})-f(r_{t-1}-r_t)+\psi}{2\psi}>0 \tag{16}$$

$$\Leftrightarrow \Big(\frac{-1-f'(r_{t-1}-1)}{2\psi}\Big)2+1-\frac{\psi}{2\psi}>0 \tag{17}$$

$$\Leftrightarrow \frac{-1-f'(r_{t-1}-1)}{\psi}+\frac{1}{2}>0 \tag{18}$$

$$\Leftrightarrow \frac{\psi}{2(1+f'(r_{t-1}-1))}>1. \tag{19}$$

However, $\frac{\psi}{2(1+f'(r_{t-1}-1))}<\frac{\psi}{2}$ and, by assumption, $\frac{\psi}{2}<1$, a contradiction. Thus, there cannot be an equilibrium $r^*=1$.

Now consider an interior equilibrium, $r_t^* \in (0,1)$. Substituting $r_t=-d_t$ into R's first-order condition (12), we obtain the following expression:

$$2r_t\Big(-1 - f'(r_{t-1} - r_t)\Big) + \psi = 0. \tag{20}$$

Thus, if r_t^* is a symmetric equilibrium, then it must solve

$$r_t = \frac{\psi}{2\Big(1 + f'(\Delta r_t)\Big)}. \tag{21}$$

Finally, we prove that symmetric equilibrium exists and is unique. Define the function $h : [0,1] \to [-1,1]$ as

$$h(r) = \frac{\psi}{2\Big(1 + f'(r_{t-1} - r)\Big)} - r.$$

By equation (21), if r_t^* is a symmetric equilibrium then $h(r_t^*) = 0$. We first show that there exists a solution to the equation $h(r) = 0$.

At $r = 0$ we have:

$$h(0) = \frac{\psi}{2\Big(1 + f'(r_{t-1})\Big)} > 0.$$

Instead, at $r = 1$:

$$h(1) = \frac{\psi}{2(1 + f'(r_{t-1} - 1))} - 1 < \frac{\psi}{2} - 1 < 0.$$

Furthermore, continuity of $f'(\Delta r_t)$ implies continuity of h. Thus, by the intermediate value theorem, there exists an r_t^* such that $h(r_t^*) = 0$.

Next, we demonstrate that this solution is unique. Differentiating yields

$$h'(r) = \frac{\psi}{2}\frac{f''(r_{t-1} - r)}{1 + f'(r_{t-1} - r)} - 1.$$

By Assumption 1, $\frac{\psi}{2} < 1$ and $\frac{f''(r_{t-1}-r)}{1+f'(r_{t-1}-r)} < 1$. Hence, $h'(r) < 0$ and there is only one r_t^* such that $h(r_t^*) = 0$.

Finally, we show that r_t^* is indeed R's best response to D choosing $-r_t^*$. Consider R's first-order condition with D's policy fixed at $-r_t^*$:

$$\Big(\frac{-1 + f'(r_{t-1} - r_t)}{2\psi}\Big)(r_t + r_t^*) + 1 - \frac{r_t - r_t^* + f(-r_t^* - d_{t-1}) - f(r_{t-1} - r_t) + \psi}{2\psi} = 0. \tag{22}$$

By construction r_t^* solves R's first-order condition. To conclude the argument note that, by Lemma A3, R's expected utility is concave in r_t, thus, r_t^* must be the unique solution to R's first-order condition. The symmetric argument proves that $d_t^* = -r_t^*$ is D's best response to

R choosing r_t^*, as required. ■

Proof of Corollary 2. The result follows from Proposition 1, Corollary 1, and the observation that $f'(\cdot) > 0$ everywhere. ■

Lemma A4 *If* $r_0 = \overline{r}$ *then* $r_1^* = \overline{r}$. *If* $r_0 < \overline{r}$ *then* $r_1^* > \overline{r}$. *Otherwise, if* $r_0 > \overline{r}$ *then* $r_1^* \in (\overline{r}, r_0)$.

Proof. Recall that r_1^* must solve $h(r) = 0$, with h defined as in the proof of Proposition 1. First, let $r_0 < \overline{r}$. We show that $h(r) > 0$ for all $r < \overline{r}$ and, hence, it must be that $r_1^* \geq \overline{r}$. Notice that $\frac{\psi}{2(1+f'(r_0-r))}$ is minimized at $r = r_0$ because f is s-shaped. Therefore, for $r < \overline{r}$, $h(r) > \frac{\psi}{2(1+f'(0))} - r = \overline{r} - r \geq 0$, as required.

Finally, suppose $r_0 > \overline{r}$. We demonstrate that if r_1^* solves $h(r) = 0$ then $r_1^* \in (\overline{r}, r_0)$. Evaluating h at $r = r_0$ yields $h(r_0) = \frac{\psi}{2(1+f'(0))} - r_0 = \overline{r} - r_0 < 0$. Instead, at $r = \overline{r}$ we have $h(\overline{r}) = \frac{\psi}{2(1+f'(r_0-\overline{r}))} - \overline{r} > 0$, where the final inequality follows because $\overline{r} = \frac{\psi}{2(1+f'(0))}$ and $\frac{\psi}{2(1+f'(r_0-r))}$ is minimized at $r = r_0$ when f is s-shaped. Thus, $r_1^* \in (r_0, \overline{r})$.

To conclude, note that r_1^* is continuous in r_0 and, thus, if $r_0 = \overline{r}$ then $r_1^* = \overline{r}$. ■

Lemma A5

1. *If* $r_0 < \overline{r}$ *then* r_1^* *is decreasing in* r_0. *Otherwise, if* $r_0 > \overline{r}$ *then* r_1^* *is increasing in* r_0.
2. $\left|\frac{\partial r_1^*}{\partial r_0}\right| < 1$.

Proof. Applying the implicit function theorem to (8):

$$\frac{\partial r_1^*}{\partial r_0} = \frac{\frac{-\psi f''(r_0-r_1^*)}{2(1+f'(r_0-r_1^*))^2}}{1 - \frac{\psi f''(r_0-r_1^*)}{2(1+f'(r_0-r_1^*))^2}} \tag{23}$$

By Assumption 1, $\frac{\psi}{2} < 1$ and $\frac{f''(r_0-r_1^*)}{1+f'(r_0-r_1^*)} < 1$. Thus, the denominator of (23) is positive. Therefore, $\frac{\partial r_1^*}{\partial r_0} > 0$ if and only if $f''(r_0 - r_1^*) < 0$.

By Lemma A4, $r_1^* > r_0$ if and only if $r_0 < \overline{r}$. Since f is s-shaped, this implies $f''(r_0 - r_1^*) > 0$ for $r_0 < \overline{r}$ and $f''(r_0 - r^*) < 0$ for $r_0 > \overline{r}$. Thus, $\frac{\partial r_1^*}{\partial r_0} < 0$ if $r_0 < \overline{r}$ and $\frac{\partial r_1^*}{\partial r_0} > 0$ if $r_0 > \overline{r}$, as claimed.

To conclude, we show that $|\frac{\partial r_1^*}{\partial r_0}| < 1$. Clearly $\frac{\partial r_1^*}{\partial r_0} < 1$. Finally, by Assumption 1, $\frac{\psi}{2} < 1 < 1 + f'(r_0 - r_1^*)$ and $f''(r_0 - r_1^*) < 1 + f'(r_0 - r_1^*)$, which implies $\frac{\partial r_1^*}{\partial r_0} > -1$. ■

Proof of Proposition 2. Follow immediately from Lemmas A4 and A5. ■

Proof of Proposition 3. To show part 1, first note that, by Proposition 2, if $r^*_{\bar{t}-1} > \overline{r}$ at period $\bar{t}$, then $r^*_{\bar{t}} \in (\overline{r}, r^*_{\bar{t}-1})$. Thus, the sequence of platforms $(r^*_{\bar{t}-1}, r^*_{\bar{t}}, r^*_{\bar{t}+1}, \dots)$ is strictly decreasing and bounded below by $\overline{r}$.

We now show that $\overline{r}$ is the infimum of $(r^*_{\bar{t}-1}, r^*_{\bar{t}}, r^*_{\bar{t}+1}, \dots)$. Suppose instead that the infimum is given by some $r' > \overline{r}$. Then the monotone convergence theorem implies that $\lim_{t\to\infty} r^*_t = r'$. However, by Proposition 2 if $r^*_t = r' > \overline{r}$ then $r^*_{t+1} < r'$. Therefore, for all $\epsilon > 0$ sufficiently small if $r^*_t = r' - \epsilon$ then $r^*_{t+1} < r'$, contradicting that $r' > \overline{r}$ is the infimum of the sequence of platforms.

Thus, $\overline{r}$ is the infimum of $(r^*_{\bar{t}-1}, r^*_{\bar{t}}, r^*_{\bar{t}+1}, \dots)$ and the monotone convergence theorem yields $\lim_{t\to\infty} r^*_t = \overline{r}$.

Parts 2 and 3 then follow immediately from Proposition 2. ■

Proof of Proposition 4. The proof for when the state does not change, $\psi_t = \psi_{t-1}$, follows immediately by Proposition 3. Next, suppose that $\psi_{t-1} = \psi_h$ and $\psi_t = \psi_l$. First, by Proposition 3, $r^*_t \in (\overline{r}_l, r^*_{t-1})$ for any r^*_{t-1}, as claimed. Second, note that if $\psi_{t-1} = \psi_h$ then $r^*_{t-1} > \overline{r}_h > \overline{r}_l$ by Proposition 2. Thus, Proposition 2 also yields $\frac{\partial r^*_t}{\partial r_{t-1}} < 1$, which implies that $r^*_t - r^*_{t-1}$ is increasing in r^*_{t-1}.

Similarly, if the state changes from $\psi_{t-1} = \psi_l$ to $\psi_t = \psi_h$, and $r^*_{t-1} > \overline{r}_h$, Proposition 2 implies that $r^*_{t-1} > r^*_t > \overline{r}_h$ and $r^*_{t-1} - r^*_t$ is increasing in r^*_{t-1}.

Finally, if the state changes from $\psi_{t-1} = \psi_l$ to $\psi_t = \psi_h$, and $r^*_{t-1} < \overline{r}_h$, Proposition 2 immediately implies that $r^*_t > \overline{r}_h > r^*_{t-1}$, and $r^*_t - r^*_{t-1}$ is decreasing in r^*_{t-1}. ■

The change in the state ψ_t together with the equilibrium platform r^*_t forms a Markov chain on $[0,1] \times \{\psi_l, \psi_h\}$. Abusing notation, for $\psi_j \in \{\psi_l, \psi_h\}$ let $\psi_{-j} \in \{\psi_l, \psi_h\}$ be the state not equal to ψ_j. Additionally, define $r^*_j(r)$ as the equilibrium policy when the state is ψ_j and the anchor is r. Then, for given state $(\psi_t, r_t) = (\psi, r)$ the transition kernel can be described by: (1) Draw $\psi_{t+1} \in \{\psi_l, \psi_h\}$ according to $Pr(\psi_{t+1} = \psi_j | \psi_t = \psi_j, r_t = r) = \rho$ and $Pr(\psi_{t+1} = \psi_{-j} | \psi_{t+1} = \psi_j, t_t = r) = 1 - \rho$; and (2) For $\psi_{t+1} = \psi_{j'}$ set $r_{t+1} = r^*_{j'}(r)$. Thus, from any current state (r, ψ_j), the chain can only move to one of two possible next states: $(r^*_l(r), \psi_l)$ or $(r^*_h(r), \psi_h)$, with probabilities ρ and $1 - \rho$ depending on the current state.

Before proving the results in the text notice that our Markov chain has a continuous state space, and thus we must show that a stationary distribution actually exists. A stationary distribution of this Markov chain will be given by a probability measure π on $[0,1] \times \{\psi_l, \psi_h\}$. Define $\pi_\psi = (\pi_l, \pi_h)$ as the marginal distribution of π over $\{\psi_l, \psi_h\}$. We first prove that in the stationary distribution we must have $\pi_l = \pi_h = 1/2$.

Lemma A6 $\pi_l = \pi_h = 1/2$.

Proof. The draw of $\psi_t \in \{\psi_l, \psi_h\}$ is independent of r_t and given by $Pr(\psi_{t+1} = \psi_t) = \rho$. Thus, π_ψ solves:

$$\pi_l = \rho\pi_l + (1-\rho)\pi_h$$
$$\pi_h = (1-\rho)\pi_l + \rho\pi_h.$$

This implies $\pi_l = \pi_h$ and together with the requirement that $\pi_l + \pi_h = 1$ this delivers $\pi_l = \pi_h = 1/2$. ■

For a Borel measurable $A \subseteq [0,1]$ let $\mu_j(A) = Pr(r_t \in A|\psi_t = \psi_j)$ be the distribution of π on A conditional on $\psi = \psi_j$. Together with Lemma A6, π is then fully determined by the conditional distributions μ_l and μ_h. Hence, the stationary distribution π is characterized by a pair of measures (μ_l, μ_h) that satisfy:

$$\mu_l(A) = \rho\mu_l\big(\{r \in [0,1] | r_l^*(r) \in A\}\big) + (1-\rho)\mu_h\big(\{r \in [0,1] | r_l^*(r) \in A\}\big) \tag{24}$$
$$\mu_h(A) = \rho\mu_h\big(\{r \in [0,1] | r_h^*(r) \in A\}\big) + (1-\rho)\mu_l\big(\{r \in [0,1] | r_h^*(r) \in A\}\big) \tag{25}$$

Before proving that a solution exists and is unique, we state a helpful lemma.

Lemma A7 *For $j \in \{l, h\}$ there exists $\kappa_j < 1$ such that $|r_j^*(x) - r_j^*(y)| \leq \kappa_j |x - y|$ for all $x, y \in [0,1]$.*

Proof. Follows from part 3 of Lemma A5 and the mean value theorem. ■

Proposition 6 *There exists a unique stationary distribution.*

Proof. Let $\mathcal{P}([0,1])$ be the set of Borel probability measures on $[0,1]$ equipped with the 1-Wasserstein distance W_1.

Define the operator $T : \mathcal{P}([0,1])^2 \to \mathcal{P}([0,1])^2$ as:

$$T(\mu_l, \mu_h) = \Big(T_l(\mu_l, \mu_h), T_h(\mu_l, \mu_h)\Big),$$

where

$$T_l(\mu_l, \mu_h)(A) = r_l^* \# \big(\rho\mu_l + (1-\rho)\mu_h\big)(A)$$
$$T_h(\mu_l, \mu_h)(A) = r_h^* \# \big((1-\rho)\mu_l + \rho\mu_h\big)(A).$$

The notation $r_j^*\#(\cdot)(A)$ denotes the pushforward of the mixture measure through r_j^* for a Borel measurable set $A \subseteq [0,1]$. That is,

$$r_l^*\#(\rho\mu_l + (1-\rho)\mu_h)(A) = \rho\mu_l\big(r \in [0,1] : r_l^*(r) \in A\big) + (1-\rho)\mu_h\big(r \in [0,1] : r_l^*(r) \in A\big)$$
$$r_h^*\#((1-\rho)\mu_l + \rho\mu_h)(A) = (1-\rho)\mu_l\big(r \in [0,1] : r_h^*(r) \in A\big) + \rho\mu_h\big(r \in [0,1] : r_h^*(r) \in A\big)$$

Therefore, a fixed point of T corresponds to a solution of (24) and (25).

To show that a fixed point exists and is unique we argue that T is a contraction mapping. We prove there exists $\kappa' < 1$ such that for all (μ_l, μ_h) and $(\tilde{\mu}_l, \tilde{\mu}_h)$ in $\mathcal{P}([0,1])^2$:

$$d\Big(T(\mu_l, \mu_h), T(\tilde{\mu}_l, \tilde{\mu}_h)\Big) \leq \kappa' d\Big((\mu_l, \mu_h), (\tilde{\mu}_l, \tilde{\mu}_h)\Big). \tag{26}$$

With the sup metric on the product space $\mathcal{P}([0,1])^2$ we can write (26) as:

$$\begin{aligned}
&\max\Big\{W_1\big(r_l^*\#(\rho\mu_l + (1-\rho)\mu_h), r_l^*\#(\rho\tilde{\mu}_l + (1-\rho)\tilde{\mu}_h)\big),\\
&W_1\big(r_h^*\#((1-\rho)\mu_l + \rho\mu_h), r_h^*\#((1-\rho)\tilde{\mu}_l + \rho\tilde{\mu}_h)\big)\Big\}\\
&\leq \kappa' \max\Big\{W_1\big(\rho\mu_l + (1-\rho)\mu_h, \rho\tilde{\mu}_l + (1-\rho)\tilde{\mu}_h\big), W_1\big((1-\rho)\mu_l + \rho\mu_h, (1-\rho)\tilde{\mu}_l + \rho\tilde{\mu}_h\big)\Big\}
\end{aligned}$$

By Lemma A7 we know that r_j^* is a contraction with Lipschitz constant κ_j, hence, $W_1(r_j^*\#\lambda, r_j^*\#\tilde{\lambda}) \leq \kappa_j W_1(\lambda, \tilde{\lambda})$ (see, e.g., Villani, 2008). Therefore,

$$\begin{aligned}
&\max\Big\{W_1\big(r_l^*\#(\rho\mu_l + (1-\rho)\mu_h), r_l^*\#(\rho\tilde{\mu}_l + (1-\rho)\tilde{\mu}_h)\big),\\
&\quad W_1\big(r_h^*\#((1-\rho)\mu_l + \rho\mu_h), r_h^*\#((1-\rho)\tilde{\mu}_l + \rho\tilde{\mu}_h)\big)\Big\}\\
&\quad \leq \max\Big\{\kappa_l W_1\big(\rho\mu_l + (1-\rho)\mu_h, \rho\tilde{\mu}_l + (1-\rho)\tilde{\mu}_h\big), \kappa_h W_1\big((1-\rho)\mu_l + \rho\mu_h, (1-\rho)\tilde{\mu}_l + \rho\tilde{\mu}_h\big)\Big\},
\end{aligned}$$

and we obtain the result by defining $\kappa' = \max\{\kappa_l, \kappa_h\} < 1$.

Because T is a contraction and $\mathcal{P}([0,1])^2$ is a complete metric space, by the contraction mapping theorem there exists a unique fixed point (μ_l^*, μ_h^*), and thus a unique stationary distribution. ■

Proof of Proposition 5. First, we show that the probability that we observe platforms becoming more extreme is strictly less than $\frac{1}{2}$, which proves the first bullet point.

Recall that platforms become more extreme if and only if, at time t, the state changes from ψ_l to ψ_h *and* the anchor is below $\bar{r}_h$. Furthermore, by Lemma A6 $\Pr(\psi_t = \psi_h) = \Pr(\psi_t = \psi_l) = \frac{1}{2}$. Thus, we have that $\Pr(x_t > x_{t-1}) = \Pr(\psi_{t-1} = \psi_l)\Pr(\psi_t = \psi_h|\psi_{t-1} = \psi_l) \times \Pr(x_{t-1} < \bar{r}_h) \leq \Pr(\psi_{t-1} = \psi_l)\Pr(\psi_t = \psi_h|\psi_{t-1} = \psi_l) = \frac{1}{2}(1-\rho) < \frac{1}{2}$.

Next, consider the second bullet point. Recall that, fixing the state, the platform is U-shaped in the anchor. In the stationary distribution, the most extreme platform emerges when the anchor is at $\overline{r}_l$, and the state is ψ_h. Denote this platform x^e, and let $\overline{\delta}_{out} \equiv x^e - \overline{r}_l$ be the largest outward movement we can sustain in equilibrium. Furthermore, because the size of inward movement is decreasing in the anchor, this also implies that the largest jump inward occurs when the anchor is at x^e and the state is ψ_l. Denote the resulting platform $x(x^e, \psi_l) < \overline{r}_h$, and let $\overline{\delta}_{in} \equiv x^e - x(x^e, \psi_l)$. Notice that, by definition, $\overline{\delta}_{out} > \overline{\delta}_{in}$.

It is straightforward to see that in a stationary distribution we must have $\Pr(x_t - x_{t-1} > \overline{\delta}_{in}) > 0 = \Pr(x_{t-1} - x_t > \overline{\delta}_{in})$. The last equality follows from the definition of $\overline{\delta}_{in}$. Consider now the first inequality. Given Proposition 4, continuity implies that there exists a unique $x^\dagger \in \left(\overline{r}_l, x(x^e, \psi_l)\right)$ such that, if $x_{t-1} \in \left(\overline{r}_l, x^\dagger\right)$ and $\psi_t = \psi_h$, then $x_t - x_{t-1} > \overline{\delta}_{in}$. Thus, $\Pr(x_t - x_{t-1} > \overline{\delta}_{in}) = \Pr\left(x_{t-1} \in (\overline{r}_l, x^\dagger)\right) \times \Pr(\psi_t = \psi_h)$. In the stationary distribution, $\Pr(\psi_t = \psi_h) = \frac{1}{2}$. Furthermore, $\Pr\left(x_{t-1} \in (\overline{r}_l, x^\dagger)\right) > 0$, since this will occur in any path where $\psi_{t'} = \psi_l$, the state remains unchanged until time $t > t'$, and t is sufficiently large. ■