Argumentation Strategies in Party Competition

Catherine Hafer Federica Izzo Dimitri Landa

February 4, 2025

Abstract

Political parties' rhetorical strategies play a crucial role in shaping public opinion and electoral outcomes. To gain insight into what kind of arguments parties present to the public, and under what conditions, we develop a model of argumentation where parties compete to persuade voters before engaging in platform competition. Our model allows us to explore when parties present arguments that highlight the strengths of their ideological positions, as opposed to those that expose the weaknesses of their opponents'; when parties try to persuade voters on the same or on different dimensions; and when parties tacitly collude on an dimension, neither truly attempting to change voters' preferences.

Introduction

Political parties are "opinion-forming agencies of great importance" (Campbell et al. 1960, p. 128). They continually disseminate rhetorical messages to voters through interviews, televised congressional debates, public speeches and social media posts. These rhetorical strategies can play a crucial role in shaping public opinion and, consequently, electoral outcomes (Nelson, 2004; Druckman, Fein and Leeper, 2012).

Deciding what kind of rhetorical arguments to present to the public thus becomes a strategic imperative. A widely shared perspective is that, in order to sway public opinion, parties need to "identify—and then emphasize—those considerations that work to their advantage," (Jerit, 2008 p. 2). This includes, of course, deciding which dimensions to prioritize when persuading voters, what kind of arguments to articulate, and whether to attempt to influence the salience voters attribute to various dimensions or, instead, their ideological preferences.

The political science literature has paid little attention to providing a theoretical framework for understanding these choices. We aim to address this gap by introducing a game-theoretic model of rhetorical argumentation within a spatial elections framework, where parties compete to persuade voters and then set their electoral platforms. While existing works on electoral competition usually focus on how parties design platforms to cater to voters' exogenous preferences, our contribution complements this literature by allowing parties to influence voters' preferences and shape the political environment *before* engaging in platform competition.

Rhetoric and Persuasion: Our Approach

We consider two parties competing for the support of a voter on multiple dimensions. The model has two stages: a persuasion stage, in which the parties present rhetorical arguments, and an electoral stage, in which they set platforms. The players face common uncertainty about which dimensions are relevant for the voter and the location of her optimal platform on these dimensions. For example, consider a middle-class voter deciding whether to support a redistributive policy.

While the voter may not benefit directly, redistribution could offer indirect advantages, such as boosting consumer spending and economic growth. Conversely, it might cause harm if higher taxes reduce investment and hurt the economy. Alternatively, both effects might be negligible, making the dimension irrelevant to the voter's welfare. Thus, on each dimension, the voter can be one of three types: left-wing, right-wing, or unconcerned (finding the dimension to be broadly irrelevant or, alternatively, finding that all policy alternatives on that dimension are equally good or equally bad).[1]

While the voter's unknown type describes her innate preferences, parties' rhetorical arguments may influence her beliefs about her type and thus her induced preferences over policies. On each dimension, parties choose whether to present *supporting* arguments that aim to convince the voter that a policy aligned with their own ideology is the best choice for her, *refuting* arguments that aim to discredit policies aligned with the opponent's, or *vacuous* arguments that do not meaningfully engage with the dimension at hand.

As an example of politicians using refuting arguments, consider a recent speech by Elizabeth Warren on redistribution.[2] Warren's rhetoric focused on a critique of trickle-down economics: "trickle-down just means helping the biggest corporations and the richest people in this country, and claiming that those big corporations and rich people could be counted to create an economy that would work for everyone else." These claims, she goes on to argue, "never really made much sense ... The top 10% got all the growth in income over the past 30 years–all of it–and the economy stopped working for everyone else." This example captures the core of refuting arguments in our model: Warren's rhetorical approach focused on highlighting the weaknesses and flaws in arguments commonly used to defend conservative economic policies, rather than presenting reasons to convince

[1]Our approach does not assume that all voters must face this type of uncertainty over all policy dimensions. However, for persuasion to occur, at least *some* voters must be uncertain about the optimal policy on at least *some* dimensions. Our focus is on these voters and dimensions.

[2]https://www.warren.senate.gov/newsroom/press-releases/senator-warren-and-039s-remarks-at-afl-cio-national-summit-on-raising-wages

voters of the benefits of her own redistributive platforms.

Contrast this with a statement by former U.S. President Barack Obama,[3] exemplifying the use of supporting arguments: "You don't have to take a vow of poverty just to say, 'Well, let me help out... let me look at that child out there who doesn't have enough to eat or needs some school fees, let me help him out. I'll pay a little more in taxes... When economic power is concentrated in the hands of the few, history also shows that political power is sure to follow — and that dynamic eats away at democracy." Rather than criticizing trickle-down economic theories, Obama presents an argument directly *in support* of redistribution, by emphasizing moral aspects of the issue as well as the importance of reducing inequalities for democratic stability.

In modeling the effect of parties' arguments, whether supporting or refuting, we depart from the cheap-talk, asymmetric-information framework typically used to model verbal persuasion. Cheap-talk emphasizes the speaker's credibility: the speaker possesses information that is unknown to the receiver, and persuasion is successful when the receiver believes that the speaker is not misrepresenting this information.[4] This framework is valuable for understanding the strategic effect of expertise, but "cannot make sense of the internal persuasive force" of argumentation (Minozzi and Siegel 2010, p. 7).

In contrast, we think about a setting where politicians do not have private information about which policy is best for the voter and their arguments aim to invoke knowledge she already possesses, encouraging her to apply this knowledge to the issue at hand to draw the intended logical, factual, or normative conclusions (Hafer and Landa, 2007). Arguments that convince do so because they "make sense" to the voter; arguments that fail to convince inadvertently expose weaknesses in the party's case (Wood and Porter 2019 p. 141).

This approach aligns with an important strand of the empirical scholarship, which has empha-

[3]https://www.cnbc.com/2018/07/18/barack-obama-on-wealth-inequality-only-so-much-you-can-eat.html

[4]A related approach based on private information with verifiable disclosure assumes that messages are always persuasive when they are presented (see, e.g., Dziuda, 2011).

sized that voters are not merely passive recipients of political parties' persuasive efforts (Chong and Druckman, 2007), but, rather, engage in a process of "deliberate integration," evaluating the relevance and significance of ideas presented by the parties (Nelson, Clawson and Oxley 1997, p. 578). Because our focus is on the parties' strategic incentives, we represent the process by which voters receive arguments and adjust their preferences simply. We explicitly model this process as Bayesian updating, however our framework can be interpreted more broadly: parties may be thought to successfully (or unsuccessfully) persuade the voter either because she "impartially" evaluates the strength of their arguments or because she is a type that is inherently more (or less) susceptible to arguments with a particular ideological connotation. In the latter interpretation, effective persuasion may critically depend on the voter's emotional reactions, including a "gut feeling" that certain themes or considerations are relevant (Nabi, 1999; Wirz, 2018), and on the role of emotions in facilitating internal deliberation by making relevant memories and analogies more accessible (Marcus et al., 2005; Kühne et al., 2011).

Formally, the voter's type dictates both which arguments resonate with her *and* her optimal policy on each dimension. The types of voters who are more easily swayed by left-wing (right-wing) arguments are also the ones who tend to prefer left-wing (right-wing) policies. When an argument resonates, it moves the voter's beliefs (and induced preferences) in the speaker's optimal direction; when it doesn't, it moves the voter in the opposite direction. The possibility of arguments not only failing to persuade, but ultimately leading the receiver to update *against* the speaker, is in line with the "backfiring effect" documented in the empirical literature (see, e.g., Dickson, Hafer and Landa 2008.Bail et al. 2018. Slothuus and De Vreese 2010).

In our framework, refuting and supporting arguments have two subtle but critical differences in their effects. To understand these differences, consider an unconcerned voter type (i.e., a voter who would find supporting arguments for both left-wing and right-wing policies unpersuasive). Suppose this voter receives only the left-wing party's refutation of the right-wing policy. The voter finds this refuting argument persuasive for the same reason she would find a supporting argument for the right-wing policy unpersuasive, and shifts her induced policy preference leftward. If she receives only a supporting argument from the left-wing party, she finds it unpersuasive and shifts rightward. Only the refuting argument allows the party to shift the unconcerned voter in its direction, thus refuting arguments are more effective on the *extensive margin* of persuasion (corresponding to moving the voter's induced ideal point in the party's preferred direction).

Relying exclusively on the refuting argument may come at a cost, however. Consider the effects of supporting and refuting arguments on, for example, a voter whose underlying type is left-wing. This voter would always find an argument presented by the left-wing party persuasive, whether supporting or refuting, and both types of arguments would have the same directional effect on her induced preferences on the dimension. However, the effects on the voter's evaluation of the salience of the dimension would be very different. A persuasive supporting argument increases the voter's belief in the dimension's importance, while a persuasive refuting argument decreases it. In this sense, supporting arguments are more effective on the *intensive margin* (corresponding to increasing the importance of the dimension for an aligned voter).

Preview of Results

We find that, on dimensions that are ex ante likely to be important for the voter, parties tacitly collude by presenting vacuous arguments rather than attempting to change the voter's preferences. It's important to note that this does not mean parties avoid discussing these dimensions, nor that they put less emphasis on these dimensions in their rhetorical strategies. Rather, the emergence of these vacuous arguments aligns with the observation that politicians often talk "without saying anything at all,"[5] speak in "ringing generalities,"[6] evade questions, and answer without actually answering.[7] In Italy, the term for this phenomenon is *politichese*, a rhetorical style aimed precisely at not informing or explaining anything. Our model rationalizes this kind of behavior within a

[5]https://www.mic.com/articles/13722/the-politics-of-fluff-how-politicians-say-everything-without-saying-anything-at-all

[6]https://slate.com/news-and-politics/2007/06/why-do-politicians-talk-like-that.html

[7]https://www.livescience.com/14074-politicians-question-dodging-debates.html

framework where politicians strategically choose when to talk but say nothing and when to present substantive arguments in an attempt to change voters' views.

Indeed, in equilibrium, the parties do present non-vacuous arguments, aiming to persuade, but on dimensions that are ex-ante unlikely to be salient to the voter. The nature of the arguments articulated by the parties on such dimensions depends on the features of the electoral environment. When the voter's ultimate choice of party is driven by policy concerns and parties have full flexibility to cater to her induced preferences on all dimensions, parties prioritize refuting arguments over supporting arguments. This is because the intensity of the voter's preferences on each dimension is inconsequential, i.e., only the extensive margin matters.

However, supporting arguments can emerge in equilibrium when substantial frictions arise in the electoral stage – in particular, when the voter has non-policy considerations in her choice of which party to elect, or when parties have little flexibility to change their platforms on some of the dimensions. In such cases, equilibrium platforms on one dimension are functions of the platforms and voter preferences on others, and political parties can gain from increasing the relative salience of specific policy dimensions for the voter. Thus, the incentives from the intensive margin of persuasion can come to dominate those induced by the extensive margin.

In the baseline model, parties are free to attempt persuasion on all dimensions. In reality, however, it may be challenging for the voter to receive and process the parties' arguments on all dimensions at once. In extending the analysis, we thus consider a setting in which parties are constrained to attempt persuasion only on a subset of dimensions. This constraint creates a relationship between dimensions, because attempting persuasion on one necessitates foregoing it on another. When the voter ex-ante favors extreme positions on both dimensions, in equilibrium, the parties talk past each other, attempting to persuade the voter on different dimensions. This one-sided persuasion enriches our results from the baseline model, where persuasion attempts by one party are always met with competition from the other.

Related Literature

We begin with the premise that parties' rhetorical strategies can significantly influence public opinion. One perspective in the empirical literature (Lazarsfeld, Berelson and Gaudet, 1968) challenges this premise, arguing that parties' persuasive efforts have minimal effects (see also Finkel, 1993). However, other scholars present a more nuanced picture (Sides, 2006; DellaVigna and Gentzkow, 2010; Druckman, 2022). In discussing the effect of persuasion in the context of electoral campaigns, Druckman and Miller (2004) go so far as to say that 'we have moved from viewing campaigns as having minimal effects on voters to seeing them as events that can fundamentally alter election outcomes' (p. 502).

In a related debate, the "partisan intoxication" hypothesis — which suggests voters reject information inconsistent with their prior attitudes — is contrasted with a more recent perspective proposing that voters process information in a rational or quasi-Bayesian manner (Hill, 2017; Fowler et al., 2020). Both perspectives agree that partisanship (i.e., voters' initial beliefs) is a critical mediator in political persuasion, a point that aligns with our modeling assumptions. Furthermore, in line with the latter perspective, and especially relevant for our work, recent evidence suggests that persuasion can be effective even with ex-ante misaligned voters when such voters have weak prior beliefs and weak partisan attachments, and when messages have informative content (Broockman and Kalla, 2023). All these features are consistent with our model, which is therefore in line with a perspective where persuasion is possible, but hard to achieve.[8]

Building on Hafer and Landa (2007), we model a setting where speakers and receivers have symmetric information about the receivers' optimal policy and where arguments are persuasive when they resonate with the receiver's type. Unlike Hafer and Landa (2007), who focus on group deliberation and the allocation of time between presenting and listening to arguments, we embed this argumentation framework within an electoral competition model, examining parties' strategic decisions of which dimensions to address and whether to highlight their own preferred policy's

[8]In our model, a voter's ex-ante ideological leaning is a function of her prior. Weak priors correspond to more centrist initial preferences (i.e., weak partisan attachments).

strengths or their opponent's weaknesses.

Our approach also connects to the literature on Bayesian persuasion (Kamenica and Gentzkow, 2011), as presenting arguments in our framework is analogous to running a public experiment. For clarity of exposition, we defer discussion of the differences between these approaches until after the presentation of the model.

Additionally, our work engages with recent political economy literature on narratives in political persuasion (e.g., Eliaz and Spiegler 2020, Benabou and Tirole 2006, Levy, Razin and Young 2022, Izzo, Martin and Callander 2023). In the paper most directly relevant to parties' attempts at persuasion, Izzo, Martin and Callander (2023), parties present alternative models of the world, voters evaluate each model or narrative as a whole and choose the one that best explains their experiences. In contrast, in our framework, voters evaluate individual arguments based on their merit. As we elaborate further below, this allows us to focus on parties' incentives to shape public opinion *before* committing to policy positions across various dimensions.

Finally, this paper is distinct from, yet complements, the substantial body of work on electoral campaigns. Theoretical studies, such as Aragonès, Castanheira and Giani (2015) and Dragu and Fan (2016), analyze how parties choose campaign messages to influence voters' policy attitudes.[9] Empirical research explores parties' use of positive versus negative campaign strategies (e.g., Walter and Nai, 2015; Geer, 2008; Brooks and Geer, 2007; Carraro and Castelli, 2010; Lipsitz and Geer, 2017)[10] and whether campaigns involve parties talking past each other or engaging on the same dimensions (e.g., Ansolabehere and Iyengar,1994); Petrocik, 1996; Sides, 2006; Sigelman and Buell Jr, 2004, Kaplan, Park and Ridout, 2006)

Critically, these studies examine parties' behavior in the final months or weeks before elections, *after* platforms are set, whereas we study how parties select rhetorical strategies to shape voter preferences *prior* to choosing platforms. With this in mind, parties in our model observe the

[9]Less directly related, Polborn and Yi (2006) examine campaign messaging when candidates have private information about their or their opponents' quality.

[10]Our model of refuting arguments better captures policy critiques rather than personal attacks.

outcome of the argumentation stage before selecting their electoral policy positions, and are not making across-dimension aggregate arguments to support their platform. In this environment, then, parties' choice of rhetorical arguments is shaped by their ability to mitigate the inherent risks via the subsequent choice of policy platform.

Furthermore, whereas the formal works in this literature typically assume that parties can only influence the various dimensions' relative electoral salience,[11] we allow parties to influence voters' *directional* preferences.

The Baseline Model

Players and actions. We consider the strategic interaction between two policy-motivated parties, denoted L and R, and a voter, V. The parties compete in an N-dimensional policy space $\mathbb{R}^N$. Players face common uncertainty over which dimensions of the policy space are relevant to the voter and what her optimal policy is for each relevant dimension. On each dimension, the parties choose which kinds of arguments to present. Following the arguments, the voter updates her beliefs and her policy preferences. The parties then adopt binding policy platforms and the voter decides which party to elect. In this baseline model, we assume that, if the voter is indifferent between the two parties' platforms, she tosses a fair coin.

Information and payoffs. On each dimension $j \in N$, the voter could be a left-wing type, $\theta_j = -1$, a right-wing type, $\theta_j = 1$, or unconcerned $\theta_j = \emptyset$, with θ_j i.i.d. on each dimension. This independence assumption is imposed for tractability, but in Appendix D we show that our results remain unchanged if we allow the voters' types to be correlated across dimensions.

[11] Roemer (1994) allows parties to shape voters' directional preferences but assumes that when voters are exposed to a party's messages their preferences always become more aligned with the party's. Skaperdas and Grofman (1995) assumes positive campaign messages mobilize the party's supporters while negative messages demobilize the opponent's.

Formally, the voter's utility is

$$U_v = -\sum_j \mathbf{I}_j \left(\theta_j - x_j\right)^2, \tag{1}$$

where $\mathbf{I}_j = 0$ if $\theta_j = \emptyset$ and $\mathbf{I}_j = 1$ otherwise, and x_j is the implemented policy on dimension j. Note that, if dimension j is relevant for the voter (that is, $\theta_j \neq \emptyset$), her ideal policy takes value 1 or -1.

At the beginning of the game, the voter's true type is unknown to all players, including the voter herself. Players share common prior beliefs that

- $p(\theta_j = -1) = \pi_j \lambda_j$,
- $p(\theta_j = 1) = \pi_j(1 - \lambda_j)$, and
- $p(\theta_j = 0) = 1 - \pi_j$.

π_j captures the players' expectations that dimension j is salient for voter welfare, hereafter referred to as *welfare salience*.[12] λ_j is the probability that, if dimension j is welfare-salient for the voter, her ideal policy on this dimension has value -1.

Finally, for party $i \in \{L, R\}$, utility is given by

$$U_i = -\sum_j (\tilde{x}_j^i - x_j)^2, \tag{2}$$

where $\tilde{x}_j^i$ is party i's optimal policy on dimension j. For simplicity, we assume that $\tilde{x}_j^R = -\tilde{x}_j^L = 1$ for all dimensions $j \in N$. Our results would remain unchanged if we assumed that parties also obtain a benefit from winning elections per se (i.e., care about both policy and office).

Argumentation. On each dimension $j \in N$, parties simultaneously choose what kind of argument to present. In particular, each party i can present a **supporting** argument ($a_j^i = s$) that aims to convince the voter that a policy program aligned with the party's own preferences is the best

[12] We use this wording to distinguish our notion of salience from other common uses of this term in the literature.

choice for her, a **refuting** argument ($a_j^i = r$) that aims to discredit policies aligned with the opponent's preferences, or a **vacuous** argument ($a_j^i = v$) that lacks any real persuasive content. In our framework, presenting a vacuous argument is equivalent to ignoring dimension j in the rhetorical discussion and has no impact on the voter's beliefs. In contrast, non-vacuous arguments may successfully persuade the voter or backfire.

Specifically, non-vacuous arguments *resonate* with the voter if and only if their claim aligns with her underlying type θ_j: an argument supporting a left-wing (right-wing) policy resonates with the voter *if and only if* she is a left-wing (right-wing) type. In contrast, an argument refuting the left-wing (right-wing) policy resonates *unless* the voter is a left-wing (right-wing) type.

Let $\rho_{a_j^i} = 1$ denote the event that argument a_j^i resonates with the voter, and $\rho_{a_j^i} = 0$ denote the event that it does not. Table 1 summarizes our assumptions on when arguments resonate, conditional on the receiver type (we omit the subscript a_j^i for readability):

	$\theta_j = 1$	$\theta_j = -1$	$\theta_j = \emptyset$
$a_j^R = s$	$\rho = 1$	$\rho = 0$	$\rho = 0$
$a_j^R = r$	$\rho = 1$	$\rho = 0$	$\rho = 1$
$a_j^L = s$	$\rho = 0$	$\rho = 1$	$\rho = 0$
$a_j^L = r$	$\rho = 0$	$\rho = 1$	$\rho = 1$

Table 1: Argument resonance conditional on voter type and argument. a_j^R (a_j^L) denotes an argument presented by R (L).

Underlying these assumptions is the notion that each policy is associated with a set of reasons that most effectively showcase its merits. The supporting argument for a given policy highlights these reasons, while the refuting argument provides the voter with a rationale to discredit them by highlighting their ostensible logical, factual, or normative flaws. As a consequence, on each dimension, one party's supporting argument and the refuting argument from its opponent partition the type space in the same way, whereas two opposing refuting or supporting arguments do not.

Timing.

1. Parties simultaneously present N-dimensional argument vectors $\mathbf{a}^L$ and $\mathbf{a}^R$;

2. on each dimension $j \in N$, V observes arguments and their resonance and then updates beliefs on θ_j using Bayes' Rule;

3. parties observe whether arguments resonated, and simultaneously commit to platforms $\mathbf{x}^L \in \mathbf{R}^N$ and $\mathbf{x}^R \in \mathbf{R}^N$;

4. the voter chooses whom to elect;

5. the elected party implements its announced platform.

Equilibrium concept. We consider Perfect Bayesian Equilibria in pure strategies (henceforth, equilibria). When multiple equilibria exist, we focus on those that are Pareto-undominated for the parties (hereafter, Pareto-undominated equilibria).[13] The set of equilibria surviving this selection criterion exactly coincides with the Perfect Bayesian Equilibria of the variant of the game in which parties present arguments sequentially rather than simultaneously, regardless of the order or moves. In contrast, the equilibria that are Pareto-inferior are an artifact of the simultaneity of moves, as, in this game, are equilibria in mixed strategies, which occur in the baseline model only when there are multiple equilibria in pure strategies. The focus on Pareto-undominated equilibria in pure strategies should, thus, be interpreted as robustness-based.

Comments on the Baseline Model

The model prioritizes understanding parties' rhetorical choices, focusing on their strategic behavior while adopting a parsimonious approach to voter responses. This approach provides flexibility

[13]That is, we eliminate an equilibrium if, for the same parameter values, there exists another equilibrium that yields higher utility for at least one of the parties, without decreasing the utility of the other.

to explore extensions that enrich the strategic environment of the parties. Furthermore, our assumptions about voter responses to arguments are broadly consistent with empirical findings: (a) voters may accept or reject arguments (Chong and Druckman, 2007; Nelson, Clawson and Oxley, 1997); (b) arguments can backfire – shifting a voter's induced preferences away from the party's (Chong and Druckman, 2007; Nelson, Clawson and Oxley, 1997); (c) persuasion attempts are more successful with voters already more aligned with the party (Broockman and Kalla, 2023); and (d) partisan arguments that shape voters' opinions are endogenous to parties' expectations about voter receptivity to particular arguments (Zaller, 1992, 2012; Lenz, 2012), reflected in our model in the parameters λ and π, which might be derived, e.g., from opinion polls or historical observation.

The assumed responsiveness of the voter's posterior to non-vacuous arguments may be interpreted as implying that those arguments are new to the voter or, alternatively, that, even if their logic is familiar, hearing them again from the parties makes them more prominent in a voter's mind or stimulates further internal deliberation. Under the latter interpretation, arguments more familiar to the voter, or more closely related to other arguments or premises that she holds to be true, are more likely to resonate with her (and, accordingly, will have a smaller effect on her induced preferences, as her prior is already largely aligned with the corresponding positions).

The baseline setting assumes that parties have complete flexibility in their strategic choices. Of course, factors orthogonal to the incentives detailed by our model, such as pressures from the party base, may limit this flexibility. Our model elucidates parties' strategic incentives when the constraints implied by these factors do not bind but, importantly, the insights we develop remain valid even if constraints in argumentation affect different parties to different degrees (see our discussion following Lemma 2 below). Furthermore, some of the extensions we consider below study specific ways in which parties' flexibility may be limited, i.e., when parties' arguments constrain their choice of platforms and when their past decisions or reputations limit their flexibility on some dimensions.

Relatedly, parties' policy motivations discipline parties' rhetorical strategies within the constraints posed by the voter's anticipated responses and the competitive environment. (If parties were purely vote-maximizing, any set of arguments could be sustained in equilibrium in our model.)

Policy motivations also drive the trade-off between intensive and extensive margins of persuasion, highlighted in the extension with electoral frictions we analyze below.

Furthermore, while the model focuses on two-party competition and treats parties as unitary actors, our framework could be applied more broadly. In particular, (a) the model can accommodate parties as coalitions of different groups with aligned preferences, each caring predominantly about a different dimension and offering arguments on that dimension; and (b) as we detail in the Discussion section and Appendix G, the incentives we describe extend to multiparty system settings, though a fuller model of specific features of such settings would be necessary for developing more resolute predictions.

Finally, we comment on the relationship between our model of communication and learning and the Bayesian persuasion framework (Kamenica and Gentzkow, 2011). As in that framework, we can think of presenting a non-vacuous argument as conducting an experiment. Consider an analogy: imagine the voter has three cups in front of her, labeled as her three possible types, with a ball hidden under one cup indicating her true type (neither the parties nor the voter know which cup). In the context of this example, making a non-vacuous argument is equivalent to flipping one of the cups. A supporting argument flips the cup aligned with the speaker's preferred position, hoping to reveal the ball underneath, while a refuting argument flips the opponent's cup, aiming to show it is empty.[14]

However, significant differences exist between the two approaches. In the Bayesian persuasion

[14]In the baseline model, a party cannot flip multiple cups or the cup corresponding to the unconcerned voter type. We consider this richer argument space in an extension analyzed in Appendix E, where parties can choose to present a refuting, supporting, or vacuous argument, as in the baseline model; a *salience* argument, which aims to persuade the voter that she should or should not care about a specific dimension; or a combination of multiple arguments. We show that while the baseline results hold on dimensions where neither party has a strong initial advantage, salience arguments from the disadvantaged party can emerge in equilibrium when λ_j is very high or very low.

framework, the persuader can design *any* experiment, including one where, in the language of our model, an argument supporting a right-wing policy may resonate with a $\theta = -1$ voter type. In contrast, in our model, an argument can never resonate with the voter if it does not align with her type. This assumption captures the idea that arguments provide the receiver with tools to form her own opinions and beliefs.[15] Because the outcome of an argument is the result of the voter's own deliberation and reasoning, it must be consistent with her type. This suggests that our approach is particularly suitable for studying *verbal* persuasion, which works by tapping into the audience's own existing knowledge, experiences, values, or systems of beliefs (rather than by generating new evidence, as in the standard Bayesian persuasion framework).

Preliminary Analysis

Arguments and Persuasion: Single-Party Case

Before delving into equilibrium analysis, we focus on characterizing how our voter responds to ideological arguments. To fix ideas, suppose first that the voter is only interacting with one speaker - the right-wing party R. Denote the voter's preferred policy on dimension j, given her posterior beliefs, x_j^v. Table 2 displays x_j^v conditional on the resonance events indicated:

	$\rho = 1$	$\rho = 0$
$a_j^R = s$	$x_j^v = 1$	$x_j^v = -1$
$a_j^R = r$	$x_j^v = 1$	$x_j^v = -1$
$a_j^R = v$	$x_j^v = 1 - 2\lambda_j$	$x_j^v = 1 - 2\lambda_j$

Table 2: Voter's induced policy preferences, single speaker case.

[15]In contrast to the Bayesian-persuasion framework, where the information designer can choose any partition the state space, our framework limits parties to choosing whether the experiment groups the unconcerned type with the right-wing type (revealing whether $\theta = -1$) or the left-wing type (revealing whether $\theta = +1$).

To understand the effect of refuting arguments, consider a voter who finds an argument refuting the left-wing policy persuasive; in this case, left-wing policies cannot be optimal for her. However, she remains uncertain about how she would feel if confronted with an argument extolling the merits of the right-wing alternative. (Such an argument could be persuasive, indicating that policy $+1$ is optimal, or it might fail to resonate, which, combined with the resonance of the argument refuting policy -1, would indicate that policies on this dimension are inconsequential to her — i.e., she is an unconcerned type.) Because of this uncertainty regarding whether right-wing policies are optimal or merely on par with other alternatives, the voter's best choice on this dimension is policy $+1$. If, instead, the voter finds the refuting argument unpersuasive, the left-wing policy must be optimal, and she revises her preferences accordingly to $x_j^v = -1$. The voter would reason symmetrically when exposed to a supporting argument that she finds unpersuasive.

Table 2 might lead one to conclude that supporting and refuting arguments are strategically equivalent since they have the same effect on the voter's induced preferences when they resonate with her. However, and crucially for our results, these arguments differ in their likelihood of resonating, leading to distinct effects on the voter's *expected* induced preferences.

A right-wing supporting argument resonates with the voter if and only if dimension j is welfare-salient *and* her optimal policy is a right-wing one, with a probability of $\pi_j(1-\lambda_j)$. In contrast, a refuting argument by the right-wing party resonates with the voter *unless* left-wing policies are actually optimal for her, i.e., with probability $1-\pi_j\lambda_j$. Thus, the refuting argument is more likely to resonate and so yields a higher expected induced preference for the voter, $(1-\pi_j\lambda_j)-\pi_j\lambda_j > \pi_j(1-\lambda_j)-(1-\pi_j(1-\lambda_j))$.

In concluding, we note that, for the same reasons described above, refuting arguments are less effective than supporting ones in persuading the voter of the welfare salience of a particular dimension (i.e., increasing her posterior that $\theta_j \neq \emptyset$). When a supporting argument resonates, it *must* be that $\theta_j \neq \emptyset$, and the voter updates accordingly. In contrast, while a resonating refuting argument shifts the voter's preferences toward the party, it also induces her to believe that the dimension at hand is *less* likely to be welfare-salient than she previously thought. In this sense, refuting arguments are better on the extensive margin of persuasion, while supporting ones are better on the intensive margin. As we will demonstrate in a later extension, this trade-off is key to understanding parties' rhetorical strategies in the presence of electoral frictions.

Competing Rhetorical Messages

Next, we move back to our analysis of competition in persuasion, with both parties, L and R, allowed to present arguments.

	$a_j^L = s$	$a_j^L = r$	$a_j^L = v$
$a_j^R = s$	$x_j^v = \begin{cases} \mathbb{R} & \text{with probability } 1-\pi_j \\ -1 & \text{with probability } \pi_j\lambda_j \\ +1 & \text{with probability } \pi_j(1-\lambda_j) \end{cases}$	$x_j^v = \begin{cases} -1 & \text{with probability } 1-\pi_j(1-\lambda_j) \\ +1 & \text{with probability } \pi_j(1-\lambda_j) \end{cases}$	$x_j^v = \begin{cases} -1 & \text{with probability } 1-\pi_j(1-\lambda_j) \\ +1 & \text{with probability } \pi_j(1-\lambda_j) \end{cases}$
$a_j^R = r$	$x_j^v = \begin{cases} -1 & \text{with probability } \pi_j\lambda_j \\ +1 & \text{with probability } (1-\pi_j\lambda_j) \end{cases}$	$x_j^v = \begin{cases} \mathbb{R} & \text{with probability } 1-\pi_j \\ -1 & \text{with probability } \pi_j\lambda_j \\ +1 & \text{with probability } \pi_j(1-\lambda_j) \end{cases}$	$x_j^v = \begin{cases} -1 & \text{with probability } \pi_j\lambda_j \\ +1 & \text{with probability } (1-\pi_j\lambda_j) \end{cases}$
$a_j^R = v$	$x_j^v = \begin{cases} -1 & \text{with probability } \pi_j\lambda_j \\ +1 & \text{with probability } (1-\pi_j\lambda_j) \end{cases}$	$x_j^v = \begin{cases} -1 & \text{with probability } 1-\pi_j(1-\lambda_j) \\ +1 & \text{with probability } \pi_j(1-\lambda_j) \end{cases}$	$x_j^v = 1 - 2\lambda_j$

Table 3: Probability distributions over x_j^v given competition in persuasion.

If only one party presents a non-vacuous argument, voter learning follows the single-speaker case. The same applies if one party makes a supporting argument and the other a refuting one, as these two arguments partition the voter-type space in the same way. Now assume both parties present supporting arguments. With probability $\pi_j\lambda_j$, the left-wing argument resonates, leading the voter to update that dimension j is welfare-salient and the optimal policy is -1. With probability $\pi_j(1-\lambda_j)$, the right-wing argument resonates, prompting her to conclude that dimension j is welfare-salient and policy 1 is optimal. Finally, with probability $1-\pi_j$, neither argument resonates, and the voter updates that she is the unconcerned type — i.e., she becomes indifferent between policy alternatives on this dimension and will focus on other dimensions when making her electoral decision. A symmetric logic applies if both present refuting arguments.

Equilibrium Analysis: a Unidimensional Benchmark

For ease of exposition, we begin by analyzing a benchmark version of the model where the policy space consists of a single dimension. We proceed by backward induction, beginning with the platform competition stage. Denote $\widehat{\pi}$ to be the posterior probability that $\theta \neq \emptyset$. Recall that x^v denotes the policy that maximizes the voter's utility, given her posterior beliefs (as characterized in Table 3). We have:

Lemma 1. *In equilibrium, the implemented policy x^* has the following properties:*

- *if $\widehat{\pi} = 0$, then $x^* = \tilde{x}^R$ with probability $\frac{1}{2}$ and $x^* = \tilde{x}^L$ with probability $\frac{1}{2}$;*
- *if $\widehat{\pi} > 0$, then in equilibrium $x^* = x^v$.*

Parties in our model are purely policy-motivated. However, at the time of platform selection, they face no uncertainty about the voter's (induced) preferences. Consequently, the voter's preferred policy is always implemented in equilibrium when $\widehat{\pi} > 0$. To see this, conjecture an equilibrium where a policy x different from x^v is implemented with strictly positive probability. The party whose ideal point is closer to x^v than to x can always deviate closer to x^v and win with certainty while improving its payoff. Thus, no such equilibrium exists. However, if $\widehat{\pi} = 0$, the voter is indifferent to all policies and elects each party with equal probability, regardless of their platforms. In this case, each party commits to its ideal policy.

The Parties' Rhetorical Strategies

Turning to the parties' choice of rhetorical strategies, we first show that supporting arguments can never be sustained in this baseline model:

Lemma 2. *Given the equilibrium of the platform stage in the baseline model, parties always prefer making refuting arguments to making supporting arguments.*

Consider party L, and suppose R presents a vacuous argument. Then, Table 3 and Lemma 1 directly imply that L must strictly prefer a refuting argument to a supporting one, as the former is more likely to resonate and shift the voter's preferences leftward. If R presents a supporting argument, L prefers a refuting argument to a supporting one, to exploit the high probability of the opponent's argument backfiring.[16] Finally, if R presents a refuting argument, L maximizes its payoff by countering with a refuting argument to prevent the opponent from capturing the unconcerned voter type.

This result has significant implications for the robustness of the insights provided by our model. Our framework allows us to study how incentives arising from competition in persuasion shape parties' rhetorical strategies. Naturally, factors orthogonal to these incentives—such as pressures from party members or internal bargaining between factions—may also affect parties' choices in this domain, and these pressures may impact different parties in distinct ways. However, Lemma 2 demonstrates that such asymmetries would not alter the nature of the incentives arising from competition in persuasion. Even when a party is forced (for reasons external to the model) to adopt supporting arguments, its opponent will always adopt refuting arguments in response, if it has the flexibility to do so.

We now characterize the equilibrium arguments of the baseline model:

Proposition 1. *There exists a unique $\bar{\pi}(\lambda)$ such that*

- *in the unique equilibrium, both parties present refuting arguments if $\pi < \bar{\pi}(\lambda)$;*
- *in the unique Pareto-undominated equilibrium, both parties present vacuous arguments if $\pi > \bar{\pi}(\lambda)$.*

Furthermore, $\bar{\pi}(\lambda)$ increases in λ's distance from $\frac{1}{2}$.

We can always sustain an equilibrium where each party presents a refuting arguments in an attempt to undermine the attractiveness of the other's preferred policy. If one's opponent offers a refuting argument, neither a supporting nor a vacuous argument offers additional information in response. Only a refuting argument does, and its only possible effect, given the opponent's choice of

[16] Recall that a party's refuting argument is informationally equivalent to the opponent's supporting argument.

the refuting argument, is to move the unconcerned types from supporting the opponent's preferred policy to being indifferent between the parties.

However, for sufficiently high π, there is also an equilibrium in which both parties present vacuous arguments, and that equilibrium is preferred by both parties. To understand why, consider the extreme case of $\pi = 1$: if the voter receives a non-vacuous argument from either party, she learns her true position is -1 with probability λ and +1 with probability $(1 - \lambda)$ (from Table 3). If the parties were risk-neutral, they would be indifferent ex ante between this lottery and a certain policy $(1 - 2\lambda)$, but, being risk-averse, they prefer to collude tacitly on keeping the voter uninformed.

Now consider what happens to their incentives to tacitly collude on vacuous arguments as π decreases. Suppose the left-wing party deviates to a refuting argument: as before, it risks that the voter will update that left-wing policies cannot be optimal, but, the probability of this event is now only $\pi(1 - \lambda) < (1 - \lambda)$; at the same time, the probability that the voter will adopt the left-wing position has risen to $(1 - \pi) + \pi\lambda > \lambda$. If π is sufficiently small, the party's ability to capture the unconcerned type improves the lottery induced by its argumentation enough to outweigh the risk associated with it, relative to the certain policy preference of the uninformed voter. Similar logic holds for the right-wing party. Furthermore, as shown in the last part of Proposition 1, as the voter's initial attitudes become more strongly favorable towards one party, the opponent has stronger incentives to break the collusion and try to persuade the voter.[17]

Robustness

We next consider the robustness of the results to relaxing some of the most restrictive assumptions.

First, the baseline model does not account for the possibility that voters may value consistency, potentially penalizing a party for adopting a platform that contradicts its ideological arguments. For example, a party perceived as inconsistent may lose credibility, harming its immediate and

[17]This suggests that, if unmodelled (external or internal pressures) are present, disadvantaged parties' choice of arguments may be more responsive to policy concerns than would be the choice by advantaged parties.

long-term electoral prospects. We address this in an extension analyzed in Appendix F. Our analysis shows that, while the consistency constraint changes the parties' equilibrium platforms, each argument profile produces the same lottery of policies as in the baseline model. Consequently, the parties' expected payoffs from each argument profile remain unchanged, and the equilibrium of the argumentation stage is as characterized in Proposition 1.

Second, the baseline model sidesteps the issue that different types of arguments may have different likelihood of being transmitted to the voters. For instance, party communication often relies on the media, which may, e.g., favor conflict-oriented content (Schuck et al., 2017) and therefore be more likely to transmit refuting arguments than supporting ones. Alternatively, voters may be, e.g., less open to hearing refuting arguments (Ansolabehere and Iyengar, 1995).

We can can incorporate these differences without altering the qualitative insights of the baseline model. If refuting arguments are more likely to reach the voter, this asymmetry only reinforces our baseline results. If supporting arguments are more likely to be transmitted, this affects the parties' payoffs from different argument profiles but does not change their equilibrium strategies. In the baseline model, each party always prefers either a refuting or a vacuous argument to a supporting one, regardless of the opponent's strategy. Thus, even if a supporting argument is more likely to reach the voter, it is never a best response, leaving equilibrium results unchanged.

We note that when parties must rely on mediators to communicate with voters, they may be compelled to adopt simplistic rhetoric, as the media may be less inclined to transmit more complex arguments. Within the context of our model, vacuous arguments may be interpreted as capturing the most simplistic rhetoric. In this sense, one of the contributions of our model is to show that, even absent external constraints (and even absent a direct upside of simpler arguments), parties will sometimes strategically choose to make such arguments even though they are not expected to effectively shift voters' ideological preferences, in order to avoid the risk of backfiring associated with more informative rhetoric.

Multidimensional Policy Space

We next show that, absent additional changes to the environment, the findings from the benchmark model are robust to replacing a unidimensional policy space with a multidimensional one.

Proposition 2. *Consider an N-dimensional policy space. The equilibrium of the argumentation stage is as characterized in Proposition 1, with a unique cutoff $\overline{\pi}_j(\lambda_j)$ applying to each $j \in N$.*

Despite the presence of multiple dimensions, in equilibrium the parties must converge on x_j^v (as characterized in Table 3) on every dimension j that is expectationally salient for the voter ($\widehat{\pi}_j > 0$).[18] To see this, conjecture an equilibrium in which the parties do not converge on all such dimensions. Then, at least one party can move closer to x_j^v on some dimension j s.t. $\widehat{\pi}_j > 0$ and thereby increase its expected payoff by increasing probability of winning. Thus, the conjectured equilibrium cannot be sustained. In contrast, on electorally irrelevant dimensions ($\widehat{\pi}_j = 0$), the parties always propose their own ideal policies.

As in the one-dimensional case, then, the equilibrium platforms on each dimension j depend only on whether the voter assigns a positive probability to the dimension's welfare salience, and do not otherwise depend on the magnitude of the voter's posterior $\widehat{\pi}_j$. This implies that dimensions are entirely separable in equilibrium at the platform stage, with platforms on each dimension independent of those presented on the others. In our setting, this guarantees that dimensions are separable at the argumentation stage as well, and therefore, in equilibrium, parties treat each dimension as though it were the only one available. In turn, this separability implies that, the equilibrium of the argumentation stage mirrors exactly the results of the baseline model: on each dimension, the

[18] In the limiting case in which the voter learns that her type is identical across all dimensions, multiple equilibria may arise. In all these equilibria, the party that is aligned with the voter's multidimensional ideal point chooses its ideal point (which coincides with the voter's), and the other party chooses any platform. Notice, however, that all these equilibria are payoff-equivalent to the one discussed in the text, since parties do not care about platforms directly but only about the policy outcome.

equilibrium features either tacit collusion on presenting vacuous arguments or the parties competing to persuade the voter by presenting refuting arguments.

Multidimensionality with Strategic Interaction Across Dimensions

In this section, we extend our model to settings where the aforementioned separability is violated, and strategic inter-dependencies emerge across dimensions. We show that, under some conditions, this gives rise to equilibrium behavior that cannot be sustained in the baseline model.

Constrained Parties

In the baseline model, parties may present non-vacuous arguments on all policy dimensions, should they wish to do so. Here, we consider the consequences of constraining the number of dimensions on which parties may attempt to persuade voters. These constraints may result from at least two sources. The first is that parties may have limited time, media exposure or financial resources, implying that they must focus on a subset of dimensions. The second is that voters may have limited cognitive resources, which implies at once that parties may need to remind them of the logic of the arguments to significantly influence their beliefs and preferences, and that there may be a "ceiling" on the voters' capacity to actively hold arguments for a broad range of dimensions.

In this section, we assume that there are two policy dimensions, 1 and 2, but that each party is allowed to make a non-vacuous argument on one dimension at most. We impose the restriction that λ_1 and λ_2 are neither both very low nor both very high. In other words, the voter is not ex-ante too right-leaning or left-leaning on both dimensions at once. This assumption is substantively plausible given the focus of the model on two-party electoral competition between a right and a left party. Technically, it ensures the existence of pure-strategy equilibria but does not alter our qualitative insights.[19]

[19]This assumption is irrelevant in other versions of the game, as the results remain qualitatively identical for all values of λ_j. In contrast, in the extension analyzed in this section, violating the assumption on λ_1 and λ_2 leads to mixed-strategies in equilibrium. As noted above, however, mixed

As in the baseline model, in equilibrium, parties implicitly collude by presenting vacuous arguments on any dimension j s.t. $\pi_j > \overline{\pi}_j(\lambda_j)$. Thus, if π_j is sufficiently high for either dimension, the constraint analyzed in this section does not bind, and so the equilibrium is as characterized in Proposition 2. For Proposition 3, we assume $\pi_j < \bar{\pi}_j(\lambda_j)$ for all dimensions $j \in \{1, 2\}$ to focus on the cases where the constraint is relevant.

Proposition 3. *There exist unique $\widetilde{\lambda}_1(\pi_1, \pi_2)$ and $\widetilde{\lambda}_2(\pi_2, \pi_1)$, henceforth $\widetilde{\lambda}_1$ and $\widetilde{\lambda}_2$, such that*

- *there is an equilibrium in which both parties make refuting arguments on j if and only if $\lambda_{-j} \in [\widetilde{\lambda}_{-j}, 1-\widetilde{\lambda}_{-j}]$, i.e., neither party has a strong initial disadvantage on the other dimension, $-j$;*
- *there is a unique equilibrium in which the parties present refuting arguments on different dimensions if and only if for some j, $\lambda_j < \widetilde{\lambda}_j$ and $\lambda_{-j} > 1 - \widetilde{\lambda}_{-j}$. Furthermore, in this equilibrium, L makes a refuting argument on j and R on $-j$, i.e., each party makes a refuting argument on the dimension on which it has an ex ante relative disadvantage.*

In the baseline model, parties select their optimal argumentation strategies for each dimension independently, treating each dimension as if it were the only one. Conversely, the constraints analyzed in this section imply that, in equilibrium, the characteristics of one dimension influence parties' argumentation choices on the others. While this does not affect the prevalence of refuting arguments in equilibrium, it sometimes generates one-sided persuasion in equilibrium. Unlike the baseline model, where parties always counter each other's persuasion attempts, constrained parties talk past each other, focusing on different dimensions, when the voter's ex-ante beliefs are extreme on both dimesions.

To understand this result, assume that the voter initially leans heavily to the left on dimension 2 (i.e., λ_2 is high), and conjecture an equilibrium in which both parties engage on dimension 1.

strategy equilibria in our framework are an artifact of the simultaneity of moves, not robust to the possibility of sequential moves. Under sequential moves, equilibrium is always in pure strategies, and its characterization is qualitatively aligned with the results in Proposition 3.

In our framework, parties have strong incentives to counteract each other's persuasion attempt by engaging on the same dimension. However, because λ_2 is high, the right-wing party has much to gain from deviating from the conjectured equilibrium to exploit the benefits of one-sided persuasion on dimension 2. Thus, the conjectured equilibrium does not exist. A similar, symmetric rationale applies to the left-wing party when λ_2 is low. Figure 1 below provides a graphical representation of our findings in the (π_1, π_2) space.[20]

It is important to note that the implication is not that political parties avoid dimensions where they have an advantage. Rather, if the condition given above holds, the parties will discuss such dimensions by way of vacuous arguments (i.e., they will talk in *politichese* — for example, simply asserting the superiority of their preferred position or the dimension's importance). If a party is already favored by the electorate on a certain dimension, it has, given the incentives discussed so far, little reason to accept the risks involved in attempting persuasion.[21]

Electoral Frictions

We have assumed that parties have full flexibility to adopt any platform on any dimension. Consequently, parties' equilibrium platforms coincide with the voter's preferred policy (given her posterior beliefs) on any dimension she deems potentially welfare-salient (i.e., $\widehat{\pi}_j > 0$), but the relative importance the voter assigns to such dimensions is inconsequential. Thus, parties focus solely on the extensive margin of persuasion, and supporting arguments cannot be sustained in equilibrium.

[20]Recall that in the regions where $\pi_1 > \bar{\pi}_1$ and/or $\pi_2 > \bar{\pi}_2$ the constraint does not bind. In Figure 1, these are the regions above the $\pi_2 = 1 - \lambda_2$ line, and to the right of the $\pi_1 = \lambda_1$ line.

[21]As we emphasize above, our model addresses partisan argumentation in the period *prior* to electoral campaigns and studies parties' use of rhetoric to mold voters' preferences before setting their platforms. Campaigns may have a different strategic objective – e.g., to re-direct voter's attention toward dimensions where the party can exploit advantages it has consolidated at that point (see, e.g., Petrocik, 1996 and Egan, 2013 on "issue ownership") – and so may exhibit different patterns than those described in this section.

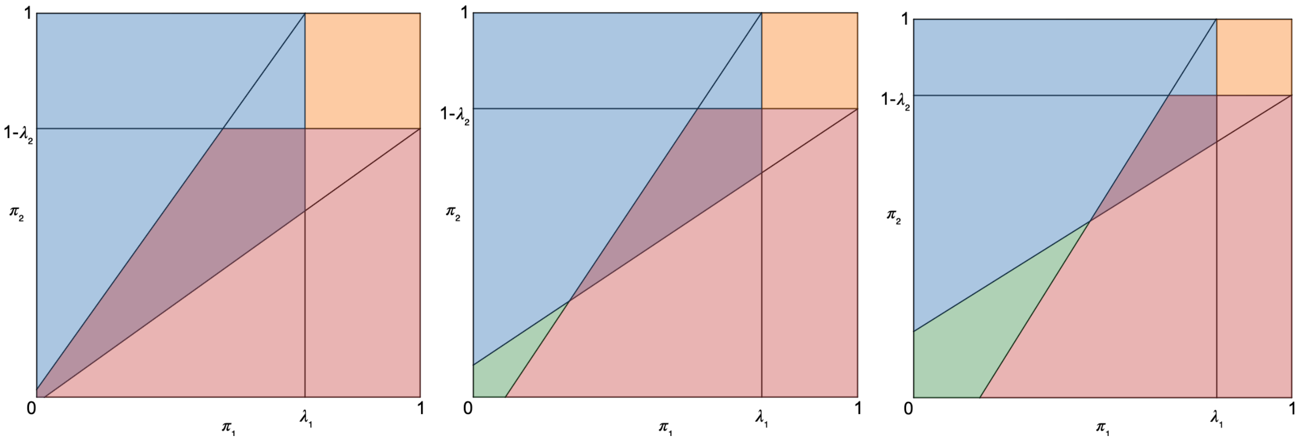

(a) $\lambda_1 = 0.7, \lambda_2 = 0.3$ (b) $\lambda_1 = 0.75, \lambda_2 = 0.25$ (c) $\lambda_1 = 0.8, \lambda_2 = 0.2$

Figure 1: Comparison of Pareto-undominated pure-strategy equilibrium behavior at different (λ_1, λ_2) with constrained parties. Red region: both parties present refuting arguments on dimension 2. Blue: both present refuting arguments on dimension 1. (purple: both present refuting on dimension 1 and both present refuting on dimension 2.) Green: L presents a refuting argument on dimension 2 and R on dimension 1. Orange: both present vacuous arguments on both dimensions.

However, frictions — limitations on parties' capacity to adapt to voters' tastes — may arise in real-world elections. For instance, voters may consider not only the parties' policy proposals, but also party leaders' fixed traits, such as experience, competence, or charisma. Alternatively, parties may be historically associated with particular positions and arguments on some of the dimensions (e.g., left-wing parties with pro-redistribution policies and right-wing parties with anti-interventionism), making shifts on these dimensions untenable on the relevant time frame and leaving each party with a fixed stance that the voter may find more or less appealing on election day. However, on newer or less ideologically entrenched dimensions, parties may have greater strategic flexibility.

In the context of our model, such frictions generate a trade-off for political parties between the intensive and extensive margins of persuasion. To study this trade-off, consider a two-dimensional issue space, consisting of a fixed dimension, on which parties take no strategic decision, and a flexible dimension, on which parties are free to choose arguments and, then, platforms. The voter's utility over the fixed dimension is determined by the realization of a shock. Formally, the voter's utility from electing party R is

$$-(x_j^v - x^{R^*})^2 + v$$

while the voter's utility from electing L is

$$-(x_j^v - x^{L^*})^2.$$

v is realized after the argumentation stage but before the platform choice, and can take one of two values: δ^R, with probability ν_R, or δ^L, with probability $1 - \nu_R$.[22] For simplicity, $\delta^R = -\delta^L = \delta > 0$. The realization of v, then, captures the net difference in leaders' popularity, or the voter's preference for one party's position on policy dimensions over which parties are historically associated with specific stances.[23]

The aim of our analysis is threefold: (1) illustrating the trade-off between extensive and intensive margins of persuasion; (2) establishing the conditional robustness of our baseline results (i.e., a prevalence of refuting arguments over supporting ones) even in the presence of this trade-off, provided electoral frictions are sufficiently small; (3) showing that sufficiently large frictions can alter the parties' strategic incentives and push them to present supporting arguments instead of refuting ones. Proposition 4 focuses precisely on establishing these results.

[22] ν_R may, in part, result from parties' rhetoric on the fixed dimension, which we take as exogenous here in order to focus on their rhetorical choices on the dimension over which they retain platform flexibility. Our findings in this regard remain robust to endogenizing ν_R.

[23] We note that the two interpretations of the model are not entirely equivalent. If the shock captures the voter's preferences over fixed policy dimensions, such dimensions should enter the parties' payoffs as well. However, if it is interpreted as capturing fixed candidates' characteristic, it should influence only the voter's utility. Importantly, this distinction does not affect the parties' strategic incentives, which depend only on the differences in utilities across argument profiles, and, therefore, as we show in Proposition 3C in Appendix C, has no impact on the results presented in this section.

Proposition 4. *There exist unique thresholds $\widehat{\delta}$, $\widetilde{\delta}$, and $\pi^{\dagger}$ such that $\widehat{\delta} \geq \widetilde{\delta}$ and*

- *an equilibrium in which both parties present supporting arguments exists and is unique if $\delta > \widehat{\delta}$ and $\pi < \pi^{\dagger}$;*
- *an equilibrium in which some party makes a supporting argument can be sustained only if $\delta > \widetilde{\delta}$.*

In this model, the party favored by the shock realization wins the election with a platform that is closer to its own ideal point than would be the winning platform in the absence of the shock. The magnitude of this difference depends on the intensity of the voter's preferences over the flexible dimension, i.e., the posterior $\widehat{\pi}$. A higher $\widehat{\pi}$ imposes more stringent constraints on the party advantaged by the shock.

To see the intuition, consider a party advantaged by the expected shock, and suppose its opponent presents a refuting argument. If the advantaged party, which expects to win, responds with a refuting argument, it will implement its ideal point when the voter is the unconcerned type, but it will make a substantial policy compromise when the voter is the type aligned with the opponent, because such a voter will have learned that the flexible policy dimension is certainly welfare-salient (i.e., $\widehat{\pi} = 1$). In contrast, if the party responds with a supporting argument, the unconcerned type and the type aligned with the opponent both learn that they share the opponent's policy position but update downward on the welfare salience of the policy dimension (i.e., $\widehat{\pi} < \pi < 1$). Consequently, the party is more likely to have to make a policy concession, but the policy concession it makes is smaller than would be the case if the voter were certain of the policy's welfare salience. Thus, the advantaged party's choice of argument amounts to optimizing the trade-off between making a larger policy concession with lower probability and making a smaller policy concession with higher probability. A symmetric logic describes the incentives that a party expecting a disadvantage from the realization of v faces when responding to a refuting argument.

Thus, in this model, parties care about both the intensive and extensive margins of persuasion, with incentives to pull the voter's ideological beliefs in their preferred direction *and* to increase the salience she places on the dimension. How parties resolve this trade-off depends on the value of δ: the larger δ, the greater the equilibrium advantage of the party favored by the shock, and the stronger the influence that $\widehat{\pi}$ has on policymaking. Thus, when δ is small, parties prefer refuting arguments over supporting ones. When δ is large, parties place greater emphasis on the intensive margin and present supporting arguments in equilibrium. When $\delta \in [\widetilde{\delta}, \widehat{\delta}]$, the game can sustain asymmetric equilibria in which one party presents a refuting argument and the other a supporting one (see Proposition 2C in Apppendix C).[24]

Proposition 4 evokes results in the empirical literature, which uncover a positive association between closeness of elections and parties' use of negative rhetoric in campaigns (Banda, 2022). In our framework, the electoral environment is ex-ante more competitive when δ is small. Proposition 4, thus, predicts that a similar pattern should also emerge in earlier stages of the electoral cycle.

Discussion

The model we introduce in this paper fills an important gap in the literature on party competition, by studying parties' use of rhetorical strategies to shape voters' policy preferences and prepare the ground for platform competition. Rather than summarizing our results (see Preview of Results section), we discuss some implications of our model that speak to future research.

Partisan argumentation in multiparty systems While our analysis has focused on competition in persuasion between two political parties, competition in multiparty settings is also of interest. Because the number and types of parties are a function of distinct features of electoral systems, and policymaking in a multiparty system is inherently more complex than in two-party systems, adequately addressing this question calls for an expansive research agenda well beyond the scope of this paper. A simple extension of our model, however, speaks to the robustness of our predictions to the possibility of multiple parties.

Consider a frictionless setting with different groups on each side of the ideological spectrum, with aligned ideologies but different issue priorities. We can accommodate any configuration of

[24]Providing further equilibrium characterization in this region would not generate additional insight.

these groups in the party system: each group may form its own party, the two groups on each side may coalesce into two main parties, or groups on one side may coalesce while the others remain fragmented. As a reduced-form approximation for a policymaking process, suppose that policymaking outcomes are monotonically influenced by the preferences of the electorate on dimensions they care about – e.g., as the electorate shifts to the right on a particular dimension, the policy outcome on that dimension tends to move rightward – and, for parsimony, let parties' payoffs be induced directly by such expectations. In equilibrium, for any configuration of the party system within this formalization, parties present vacuous arguments on dimensions for which π_j is high, and they privilege refuting arguments on other dimensions. (See Appendix G.) These results align with the findings of our baseline model.

If there is a small direct cost of argument-making, this extension also predicts "issue ownership" by particular parties on each side of the spectrum. Of course, reasons outside our model may induce parties to repeat arguments made by their ideological allies (e.g., pressures from the party's base, voters' limited recall, etc.). If such reasons are present and the costs of devising and presenting arguments are sufficiently small, the strict form of "issue ownership" identified in this version of the model may be less pronounced.

Populism and the Asymmetry of Rhetoric In many democratic regimes, an important aspect of partisan rhetoric is its potential for populist appeal. While a complete analysis of such appeal is beyond this study, our models can help elucidate the strategic incentives underlying some important aspects of populism.

A critical basis of a party's populist appeal is the party's perceived congruence with some underlying preference or identity of a subset of voters in a way that does not admit epistemic contingencies, nuance, or policy trade-offs. Such congruence may be impossible for other parties, which may be unable or unwilling to mimic those features. When the corresponding set of voters is sufficiently large, this disjunction creates a dimension on which populists have an inherent advantage. In the context of our model, the effects of the populist dimension on the voter's preferences over parties may be captured by the v parameter in the extension with electoral frictions – supposing, for this case, that the probability that the shock favors the "populist" party is high. Adopting such an interpretation of that extension, then, suggests interesting implications for parties' platforms and rhetorical choices on non-populists dimensions.

Crucially, in contrast to our baseline setting, where parties always adopt similar rhetorical strategies, competition over the populist dimension creates the possibility for asymmetries in parties' rhetoric on the non-populist dimension. Absent the populist dimension, a party always responds to a refuting argument in kind, undermining the opponent's persuasion attempt. In contrast, when competition is over both the "normal" policy and the populist dimension, we can, under some conditions, sustain an equilibrium where the populist party presents a refuting argument, while its opponent would be hurt from responding in kind and is thus pressed to present redundant (supporting or vacuous) arguments. This equilibrium illustrates how a party's advantage over the populist dimension may translate into an advantage in rhetorical competition, as it allows the populist to capture the unconcerned voter type. Finally, in our interpretation of populism, a greater ex-ante advantage of the populist party should be associated with greater platform polarization, and populist parties should be expected to present more extreme platforms than their opponents.

As the above discussion highlights, our framework is sufficiently flexible to be extended in multiple directions. By illuminating the strategic logic of partisan argumentation, we hope this paper lays the foundation for future empirical and theoretical work on how rhetorical competition shapes voter preferences, electoral outcomes, and the broader dynamics of democratic discourse.

References

Ansolabehere, Stephen and Shanto Iyengar. 1994. "Riding the wave and claiming ownership over issues: The joint effects of advertising and news coverage in campaigns." *Public opinion quarterly* 58(3):335–357.

Ansolabehere, Stephen and Shanto Iyengar. 1995. "Going negative: How political advertisements shrink and polarize the electorate.".

Aragonès, Enriqueta, Micael Castanheira and Marco Giani. 2015. "Electoral competition through issue selection." *American journal of political science* 59(1):71–90.

Bail, Christopher A, Lisa P Argyle, Taylor W Brown, John P Bumpus, Haohan Chen, MB Fallin Hunzaker, Jaemin Lee, Marcus Mann, Friedolin Merhout and Alexander Volfovsky. 2018. "Exposure to opposing views on social media can increase political polarization." *Proceedings of the National Academy of Sciences* 115(37):9216–9221.

Banda, Kevin K. 2022. "Reexamining the effects of electoral competition on negative advertising." *American Politics Research* 50(1):45–51.

Benabou, Roland and Jean Tirole. 2006. "Belief in a just world and redistributive politics." *The Quarterly journal of economics* 121(2):699–746.

Broockman, David E and Joshua L Kalla. 2023. "When and why are campaigns' persuasive effects small? Evidence from the 2020 US presidential election." *American Journal of Political Science* 67(4):833–849.

Brooks, Deborah Jordan and John G Geer. 2007. "Beyond negativity: The effects of incivility on the electorate." *American journal of political science* 51(1):1–16.

Campbell, Angus, Philip E Converse, Warren E Miller and Donald E Stokes. 1960. *The American Voter*. John Wiley & Sons.

Carraro, Luciana and Luigi Castelli. 2010. "The implicit and explicit effects of negative political campaigns: Is the source really blamed?" *Political psychology* 31(4):617–645.

Chong, Dennis and James N Druckman. 2007. "Framing public opinion in competitive democracies." *American Political Science Review* 101(4):637–655.

DellaVigna, Stefano and Matthew Gentzkow. 2010. "Persuasion: empirical evidence." *Annu. Rev. Econ.* 2(1):643–669.

Dickson, Eric S, Catherine Hafer and Dimitri Landa. 2008. "Cognition and strategy: a deliberation experiment." *The Journal of Politics* 70(4):974–989.

Dragu, Tiberiu and Xiaochen Fan. 2016. "An agenda-setting theory of electoral competition." *The Journal of Politics* 78(4):1170–1183.

Druckman, James N. 2022. "A framework for the study of persuasion." *Annual Review of Political Science* 25(1):65–88.

Druckman, James N and Joanne M Miller. 2004. "Editors' Introduction: The Political Psychology of Electoral Campaigns: Introduction to the Symposium." *Political Psychology* pp. 501–506.

Druckman, James N, Jordan Fein and Thomas J Leeper. 2012. "A source of bias in public opinion stability." *American Political Science Review* 106(2):430–454.

Dziuda, Wioletta. 2011. "Strategic argumentation." *Journal of Economic Theory* 146(4):1362–1397.

Egan, Patrick J. 2013. *Partisan priorities: How issue ownership drives and distorts American politics*. Cambridge University Press.

Eliaz, Kfir and Ran Spiegler. 2020. "A model of competing narratives." *American Economic Review* 110(12):3786–3816.

Finkel, Steven E. 1993. "Reexamining the" minimal effects" model in recent presidential campaigns." *The Journal of Politics* 55(1):1–21.

Fowler, Anthony et al. 2020. "Partisan intoxication or policy voting?" *Quarterly Journal of Political Science* 15(2):141–179.

Geer, John G. 2008. *In defense of negativity: Attack ads in presidential campaigns*. University of Chicago Press.

Hafer, Catherine and Dimitri Landa. 2007. "Deliberation as self-discovery and institutions for political speech." *Journal of theoretical Politics* 19(3):329–360.

Hill, Seth J. 2017. "Learning together slowly: Bayesian learning about political facts." *The Journal of Politics* 79(4):1403–1418.

Izzo, Federica, Gregory J Martin and Steven Callander. 2023. "Ideological Competition." *American Journal of Political Science* .

Jerit, Jennifer. 2008. "Issue framing and engagement: Rhetorical strategy in public policy debates." *Political Behavior* 30:1–24.

Kamenica, Emir and Matthew Gentzkow. 2011. "Bayesian persuasion." *American Economic Review* 101(6):2590–2615.

Kaplan, Noah, David K Park and Travis N Ridout. 2006. "Dialogue in American political campaigns? An examination of issue convergence in candidate television advertising." *American Journal of Political Science* 50(3):724–736.

Kühne, Rinaldo, Christian Schemer, Jörg Matthes and Werner Wirth. 2011. "Affective priming in political campaigns: How campaign-induced emotions prime political opinions." *International Journal of Public Opinion Research* 23(4):485–507.

Lazarsfeld, Paul F, Bernard Berelson and Hazel Gaudet. 1968. *The people's choice: How the voter makes up his mind in a presidential campaign*. Columbia University Press.

Lenz, Gabriel S. 2012. *Follow the leader?: how voters respond to politicians' policies and performance*. University of Chicago Press.

Levy, Gilat, Ronny Razin and Alwyn Young. 2022. "Misspecified Politics and the Recurrence of Populism." *American Economic Review* 112(3):928–62.
URL: *https://www.aeaweb.org/articles?id=10.1257/aer.20210154*

Lipsitz, Keena and John G Geer. 2017. "Rethinking the concept of negativity: An empirical approach." *Political Research Quarterly* 70(3):577–589.

Marcus, George E, John L Sullivan, Elizabeth Theiss-Morse and Daniel Stevens. 2005. "The emotional foundation of political cognition: The impact of extrinsic anxiety on the formation of political tolerance judgments." *Political Psychology* 26(6):949–963.

Minozzi, William and David A Siegel. 2010. A theory of deliberation as interactive reasoning. In *Annual Meeting of the Midwest Political Science Association, Chicago, Il.*

Nabi, Robin L. 1999. "A cognitive-functional model for the effects of discrete negative emotions on information processing, attitude change, and recall." *Communication theory* 9(3):292–320.

Nelson, Thomas E. 2004. "Policy goals, public rhetoric, and political attitudes." *The Journal of Politics* 66(2):581–605.

Nelson, Thomas E, Rosalee A Clawson and Zoe M Oxley. 1997. "Media framing of a civil liberties conflict and its effect on tolerance." *American Political Science Review* 91(3):567–583.

Petrocik, John R. 1996. "Issue ownership in presidential elections, with a 1980 case study." *American journal of political science* pp. 825–850.

Polborn, Mattias K and David T Yi. 2006. "Informative positive and negative campaigning." *Quarterly Journal of Political Science* 1(4):351–371.

Roemer, John E. 1994. "The strategic role of party ideology when voters are uncertain about how the economy works." *American Political Science Review* 88(2):327–335.

Schuck, Andreas RT et al. 2017. "Media malaise and political cynicism." *The international encyclopedia of media effects* pp. 1–19.

Sides, John. 2006. "The origins of campaign agendas." *British Journal of Political Science* 36(3):407–436.

Sigelman, Lee and Emmett H Buell Jr. 2004. "Avoidance or engagement? Issue convergence in US presidential campaigns, 1960–2000." *American Journal of Political Science* 48(4):650–661.

Skaperdas, Stergios and Bernard Grofman. 1995. "Modeling negative campaigning." *American Political Science Review* 89(1):49–61.

Slothuus, Rune and Claes H De Vreese. 2010. "Political parties, motivated reasoning, and issue framing effects." *The Journal of Politics* 72(3):630–645.

Walter, Annemarie S and Alessandro Nai. 2015. "Explaining the use of attack behaviour in the electoral battlefield: A literature overview." *New perspectives on negative campaigning: why attack politics matters* pp. 135–153.

Wirz, Dominique. 2018. "Persuasion through emotion? An experimental test of the emotion-eliciting nature of populist communication." *International Journal of Communication* 12:1114–1138.

Wood, Thomas and Ethan Porter. 2019. "The elusive backfire effect: Mass attitudes' steadfast factual adherence." *Political Behavior* 41:135–163.

Zaller, John. 1992. "The nature and origins of mass opinion." *Cambridge University* .

Zaller, John. 2012. "What nature and origins leaves out." *Critical Review* 24(4):569–642.

Appendix

Contents

A Proofs for Baseline Models 2

B Proofs for Multidimensionality in the Presence of Constraints 5

C Proofs for Multidimensionality with Electoral Frictions 8

D Proofs for Extension with Correlated Types 19

E Proofs for Richer Argument Space 22

F Proof for Extension with Cost of Inconsistency 26

G Proofs for Extension with Multiple Parties 28

A Proofs for Baseline Models

Lemma 1. *In equilibrium, the implemented policy x^* has the following properties:*

- *if $\widehat{\pi} = 0$, then $x^* = \tilde{x}^R$ with probability $\frac{1}{2}$ and $x^* = \tilde{x}^L$ with probability $\frac{1}{2}$;*
- *if $\widehat{\pi} > 0$, then in equilibrium $x^* = x^v$.*

Proof. If $\widehat{\pi} = 0$, the voter is indifferent between all policies and always tosses a fair coin. Thus, in equilibrium each party propose its ideal point and wins with probability $\frac{1}{2}$. Suppose instead $\widehat{\pi} > 0$. In this case, for any pair of platforms the voter elects the party proposing the platform closer to x^v, and tosses a fair coin when the platforms are equidistant from x^v. Thus, the equilibrium policy must be equal to x^v. Conjecture an equilibrium where a policy $x \neq x^v$ is implemented with strictly positive probability. The party whose ideal point is closer to x^v than to x can always deviate closer to x^v and win with certainty while improving its payoff.[1] Thus, no such equilibrium exists, and in any equilibrium x^v must be implemented with certainty when $\widehat{\pi} > 0$. □

Lemma 2. *Given the equilibrium of the platform stage in the baseline model, parties always prefer making refuting arguments to making supporting arguments.*

Proof. Denote with $\mathcal{V}^i(a^R, a^L)$ i's continuation value given arguments a^R and a^L. This continuation value incorporates expectations over the optimal platforms given the pair of argument vectors presented in the argumentation stage, and the corresponding distribution of resonance events.

$$\mathcal{V}^i(a^R = r, a^L = r) = \mathcal{V}^i(a^R = s, a^L = s) \tag{3}$$
$$= -\frac{(1-\pi)}{2}(\tilde{x}^i - \tilde{x}^{-i})^2 - \pi\lambda(-1 - \tilde{x}^i)^2 - \pi(1-\lambda)(1 - \tilde{x}^i)^2;$$

[1] If in the conjectured equilibrium the platforms are equidistant from x^v, both parties have a profitable unilateral deviation. In the conjectured equilibrium, the parties win with equal probability. Each party can win with certainty by making an arbitrarily small move in the direction of x^v. Because payoffs are continuous, this unilateral deviation is always profitable.

$$\mathcal{V}^i(a^R = r, a^L = s) = \mathcal{V}^i(a^R = r, a^L = v) = \mathcal{V}^i(a^R = v, a^L = s) \tag{4}$$
$$= -\pi\lambda(-1 - \tilde{x}^i)^2 - (1 - \pi\lambda)(1 - \tilde{x}^i)^2;$$

$$\mathcal{V}^i(a^R = s, a^L = r) = \mathcal{V}^i(a^R = s, a^L = v) = \mathcal{V}^i(a^R = v, a^L = r) \tag{5}$$
$$= -\pi(1-\lambda)(1 - \tilde{x}^i)^2 - \Big(1 - \pi(1-\lambda)\Big)(-1 - \tilde{x}^i)^2;$$

$$\mathcal{V}^i(a^R = v, a^L = v) = -(1 - 2\lambda - \tilde{x}^i)^2. \tag{6}$$

Substituting $\tilde{x}^R = -\tilde{x}^{L}_{=} = 1$, we obtain that

$$\mathcal{V}^L(a^R = s, a^L = r) = \mathcal{V}^L(a^R = s, a^L = v) = \mathcal{V}^L(a^R = v, a^L = r) \tag{7}$$
$$> \mathcal{V}^L(a^R = r, a^L = r) = \mathcal{V}^L(a^R = s, a^L = s)$$
$$> \mathcal{V}^L(a^R = r, a^L = s) = \mathcal{V}^L(a^R = r, a^L = v) = \mathcal{V}^L(a^R = v, a^L = s).$$

Similar results hold for party R. □

Proposition 1. *There exists a unique $\bar{\pi}(\lambda)$ such that*

- *in the unique equilibrium, both parties present refuting arguments if $\pi < \bar{\pi}(\lambda)$;*
- *in the unique Pareto-undominated equilibrium, both parties present vacuous arguments if $\pi > \bar{\pi}(\lambda)$.*

Furthermore, $\bar{\pi}(\lambda)$ increases in λ's distance from $\frac{1}{2}$.

Proof. The inequalities in 7 imply that we cannot sustain an equilibrium with one party presenting a refuting argument and the other a vacuous one, as the latter can profitably deviate to a refuting argument. Given Lemma 2, this only leaves two equilibrium candidates: $(a^R = v, a^L = v)$ and $(a^R = r, a^L = r)$.

First, conjecture an equilibrium in which both parties present vacuous arguments. The equilibrium exists if an only if the following conditions are jointly satisfied:

$$-\Big(\tilde{x}^L-(1-2\lambda)\Big)^2 > \max\in\Big\{-\pi\lambda\Big(\tilde{x}^L+1\Big)^2-(1-\pi\lambda)\Big(\tilde{x}^L-1\Big)^2; \\ -\pi(1-\lambda)\Big(\tilde{x}^L-1\Big)^2-\Big(1-\pi(1-\lambda)\Big)\Big(\tilde{x}^L+1\Big)^2\Big\}. \tag{8}$$

and

$$-\Big(\tilde{x}^R-(1-2\lambda)\Big)^2 > \max\in\Big\{-\pi\lambda\Big(\tilde{x}^R+1\Big)^2-(1-\pi\lambda)\Big(\tilde{x}^R-1\Big)^2; \\ -\pi(1-\lambda)\Big(\tilde{x}^R-1\Big)^2-\Big(1-\pi(1-\lambda)\Big)\Big(\tilde{x}^R+1\Big)^2\Big\}. \tag{9}$$

Setting $\tilde{x}^R = -\tilde{x}^L = 1$, rearranging and simplifying, we obtain that the equilibrium exists if and only if $\pi > \bar{\pi}(\lambda) \equiv \max\Big\{\lambda, 1-\lambda\Big\}$.

Finally, conjecture an equilibrium in which both parties present refuting arguments. Following a similar analysis as above, we can verify that the equilibrium exists if and only if $(\tilde{x}^R - x^L)^2 \leq \min\in\{2(x^L-1)^2; 2(\tilde{x}^R+1)^2\}$, which is always satisfied under $\tilde{x}^R = -\tilde{x}^L = 1$.

To conclude our proof, we apply our equilibrium selection criterion by showing that, when an equilibrium in which both parties present vacuous arguments exists, it gives both parties higher expected payoff than an equilibrium in which both present refuting arguments. This holds if and only if

$$-(1-2\lambda-\tilde{x}^i)^2 > -\frac{1-\pi}{2}(\tilde{x}^R-\tilde{x}^L)^2-\pi\lambda(-1-\tilde{x}^i)^2-\pi(1-\lambda)(1-\tilde{x}^i)^2,$$

for all $i \in \{L, R\}$. This is always satisfied under condition $\pi > \bar{\pi}(\lambda)$.

That $\bar{\pi}(\lambda)$ increases as λ moves away from $\frac{1}{2}$ follows immediately from $\bar{\pi} \equiv \max\{\lambda, 1-\lambda\}$. □

Proposition 2. *Consider an N-dimensional dimension space. The equilibrium of the argumentation stage is as characterized in Proposition 1, with a unique cutoff $\bar{\pi}_j(\lambda_j)$ applying to each $j \in N$.*

Proof. First, consider a dimension j s.t. $\widehat{\pi}_j = 0$. Such a dimension is irrelevant for the voter's electoral choice. Thus, the parties always propose their own ideal policy in equilibrium.

Second, consider a dimension j s.t. $\widehat{\pi}_j > 0$. We establish that the equilibrium policy on such dimension is always equal to x_j^v. First, assume that there exist two dimensions j' and j'' s.t. $x_{j'}^v, x_{j''}^v \in (-1, 1)$ and $x_{j'}^v \neq x_{j''}^v$. Then, in the unique equilibrium the parties converge on x_j^v. To see this, conjecture an equilibrium in which the parties do not converge on all dimensions that are expectationally relevant for the voter. Then, at least one party can move closer to x_j^v win with certainty, moving the implemented platform closer to its multidmensional ideal point. Suppose instead, $x_{j'}^v = x_{j''}^v \in \{-1, 1\}$ for all j', j''. In this limiting case, in which the voter learns that her type is identical across all dimensions, multiple equilibria may arise. In all these equilibria, the party that is aligned with the voter's multidimensional ideal point chooses its ideal point (which coincides with the voter's), and the other party chooses any platform. Notice, however, that in all these equilibria the implemented policy is equal to x_j^v on all dimensions s.t. $\widehat{\pi}_j > 0$.

Thus, the equilibrium policy on dimension j is not a function of the arguments presented on the other dimensions. Given separability of the parties' utility over dimensions, we can then express party i' expected continuation value from an argument profile $(\boldsymbol{a^R}, \boldsymbol{a^L})$, $\mathcal{V}^i(\boldsymbol{a^R}, \boldsymbol{a^L})$, as $\sum_j \mathcal{V}_j^i(a_j^R, a_j^L)$, where $\mathcal{V}_j^i(a_j^R, a_j^L)$ is as characterized in the proof of Lemma 2. Thus, the results of Proposition 1 continue to apply, with a unique cutoff $\bar{\pi}_j(\lambda_j)$ applying to each $j \in N$. □

B Proofs for Multidimensionality in the Presence of Constraints

Proposition 3.

There exist unique $\widetilde{\lambda}_1(\pi_1, \pi_2)$ *and* $\widetilde{\lambda}_2(\pi_2, \pi_1)$*, henceforth* $\widetilde{\lambda}_1$ *and* $\widetilde{\lambda}_2$*, such that*

- *there is an equilibrium in which both parties make refuting arguments on* j *if and only if* $\lambda_{-j} \in [\widetilde{\lambda}_{-j}, 1-\widetilde{\lambda}_{-j}]$*, i.e., neither party has a strong initial disadvantage on the other dimension,* $-j$*;*

- *there is a unique equilibrium in which the parties present refuting arguments on different dimensions if and only if for some* j, $\lambda_j < \widetilde{\lambda}_j$ *and* $\lambda_{-j} > 1 - \widetilde{\lambda}_{-j}$. *Furthermore, in this equilibrium,* L *makes a refuting argument on* j *and* R *on* $-j$, *i.e., each party makes a refuting argument on the dimension on which it has an ex ante relative disadvantage.*

Proof. We know from the proof of Proposition 1 that, absent a constraint, (1) refuting arguments dominate supporting ones (from the comparison of the continuation values in equations (3)-(6)); and (2) if $\pi_j > \bar{\pi}_j(\lambda_j)$, vacuous is a best response to vacuous on j, and both parties obtain higher utility from (v, v) than from (r, r). Because these results are independent of other dimensions, they hold here, too.

Suppose $\pi_j < \bar{\pi}_j(\lambda_j)$ for all j. Given the above results, the only reason why party i may choose not to present a refuting argument on j is to be able to present a refuting argument on $-j$.

Assuming that party R chooses argument r on dimension j, it is better for L to choose r on j than to choose r on dimension $-j$ iff

$$-4[\pi_j(1-\lambda_j) + \frac{1}{2}(1-\pi_j) - (1-\pi_j\lambda_j)] \geq -4[\pi_{-j}(1-\lambda_{-j}) - (1-\lambda_{-j})^2],$$

which reduces to

$$\frac{2-\pi_{-j}-\sqrt{\pi_{-j}^2+2(1-\pi_j)}}{2} \leq \lambda_{-j} \leq \frac{2-\pi_{-j}+\sqrt{\pi_{-j}^2+2(1-\pi_j)}}{2}. \tag{10}$$

Similarly, assuming that party L chooses argument r on dimension j, it is better for R to choose r on j than to choose r on dimension $-j$ iff

$$\frac{\pi_{-j}-\sqrt{\pi_{-j}^2+2(1-\pi_j)}}{2} \leq \lambda_{-j} \leq \frac{\pi_{-j}+\sqrt{\pi_{-j}^2+2(1-\pi_j)}}{2}. \tag{11}$$

Letting

$$\widetilde{\lambda}_{-j}(\pi_{-j},\pi_j) := \frac{2-\pi_{-j}-\sqrt{\pi_{-j}^2+2(1-\pi_j)}}{2} = 1-\Big(\frac{\pi_{-j}+\sqrt{\pi_{-j}^2+2(1-\pi_j)}}{2}\Big),$$

and noting that $\sqrt{\pi_{-j}^2+2(1-\pi_j)} > \pi_{-j}$ and that $\lambda_{-j} \in [0,1]$, (10) and (11 reduce to $\widetilde{\lambda}_{-j}(\pi_{-j},\pi_j) < \lambda_{-j}$ and $\lambda_{-j} < 1-\widetilde{\lambda}_{-j}(\pi_{-j},\pi_j)$, respectively.

Thus:

1. there exists an equilibrium in which both parties refute on dimension 1 iff $\widetilde{\lambda}_2(\pi_2,\pi_1) \le \lambda_2 \le 1-\widetilde{\lambda}_2(\pi_2,\pi_1)$;

2. there exists an equilibrium in which both parties refute on dimension 2 iff $\widetilde{\lambda}_1(\pi_1,\pi_2) \le \lambda_1 \le 1-\widetilde{\lambda}_1(\pi_1,\pi_2)$;

3. there exists an equilibrium in which L refutes on dimension 1 while R refutes on dimension 2 iff $\lambda_1 < \widetilde{\lambda}_1(\pi_1,\pi_2)$ and $\lambda_2 > 1-\widetilde{\lambda}_2(\pi_2,\pi_1)$;

4. there exists an equilibrium in which L refutes on dimension 2 while R refutes on dimension 1 iff $\lambda_2 < \widetilde{\lambda}_2(\pi_2,\pi_1)$ and $\lambda_1 > 1-\widetilde{\lambda}_1(\pi_1,\pi_2)$.

Note that if the conditions in the third case hold, it is not possible to satisfy the conditions for any of the other three cases, guaranteeing uniqueness. The same is true for the fourth case.

Finally, given $\lambda_j < \widetilde{\lambda}_j(\pi_j,\pi_{-j})$ and $\lambda_{-j} > 1-\widetilde{\lambda}_{-j}(\pi_{-j},\pi_j)$, $\lambda_j < \lambda_{-j}$ if $\widetilde{\lambda}_j(\pi_j,\pi_{-j}) < 1-\widetilde{\lambda}_{-j}(\pi_{-j},\pi_j)$, which is equivalent to $2 < \pi_j + \pi_{-j} + \sqrt{\pi_j^2+2(1-\pi_{-j})} + \sqrt{\pi_{-j}^2+2(1-\pi_j)}$, which holds for all $(\pi_1,\pi_2) \in [0,1]^2$. □

In what follows, let Λ_j be a mapping from the $[0,1]^2$ interval into the power set of $[0,1]$ which associates a set of λ_{-j} with a given pair (π_2,π_2). Λ_j contains all values of λ_{-j} such that, given (π_1,π_2), both parties presenting refuting arguments on j is an equilibrium. Recall that the ability to support this equilibrium does not depend on λ_j, but only on λ_{-j}.

Proposition 1B. *Let $\pi_1' > \pi_1''$. Then,*

- $\Lambda_1(\pi_1', \pi_2)$ *is a subset of* $\Lambda_1(\pi_1'', \pi_2)$*;*

- $\Lambda_2(\pi_1'', \pi_2)$ *is a subset* $\Lambda_2(\pi_1', \pi_2)$ *.*

Proof. Follows from inspection of (10) and (11). □

C Proofs for Multidimensionality with Electoral Frictions

In this section, we will first adopt the interpretation in which δ captures a valence dimension over which the parties have no preferences. Proposition 3C then shows that the results are identical if we adopt the interpretation in which δ captures the voter's (realized) preferences over a second policy dimension, on which that the parties have preferences but no flexibility over.

Proposition 4. There exist unique thresholds $\widehat{\delta}$, $\widetilde{\delta}$, and $\pi^\dagger$ such that $\widehat{\delta} \geq \widetilde{\delta}$ amd

- *an equilibrium in which both parties present supporting arguments exists and is unique if* $\delta > \widehat{\delta}$ *and* $\pi < \pi^\dagger$*;*

- *an equilibrium in which some party makes a supporting argument can be sustained only if* $\delta > \widetilde{\delta}$*.*

Proof. In order to prove Proposition 4, we will proceed with the following steps. First, we characterize the equilibrium of the platform game (Claim 1). Second, we establish the existence of the unique threshold $\widetilde{\delta}$ (Claim 2). Finally, we establish the existence of the unique pair $(\widehat{\delta}, \pi^\dagger)$ (Claim 3).

Claim 1. *If the realization of the valence shock is in L's favor, then in equilibrium* $x = \max\{-1, 1 - 2\widehat{\lambda} - \sqrt{\frac{\delta}{\widehat{\pi}}}\}$*. If, instead, the realization of the valence shock is in R's favor, then, in equilibrium* $x = \min\{+1, 1 - 2\widehat{\lambda} + \sqrt{\frac{\delta}{\widehat{\pi}}}\}$*.*

Proof. Recall that x is the policy implemented in equilibrium, while $\widehat{\lambda}$ and $\widehat{\pi}$ denote the voter's posterior beliefs on the policy dimension. If no non-vacuous argument is presented, $\widehat{\pi} = \pi$. If two supporting or refuting arguments are presented, then $\widehat{\pi}$ will either take value 0 or 1. When instead only one non-vacuous argument is presented, then $\widehat{\pi}$ will either take value 1, or an interior value strictly smaller than the prior π. In particular, applying Bayes rule, we can verify that $p(\theta \neq \emptyset | a^R = r, a^L = v, \rho_{a^R} = 1) = p(\theta = \emptyset | a^R = v, a^L = s, \rho_{a^L} = 0) = \frac{\pi(1-\lambda)}{\pi(1-\lambda)+(1-\pi)} \equiv \widehat{\pi}_l^L < \pi$. Similarly, $p(\theta \neq \emptyset | a^R = s, a^L = v, \rho_{a^R} = 0) = p(\theta = \emptyset | a^R = v, a^L = r, \rho_{a^L} = 1) = \frac{\pi\lambda}{\pi\lambda+(1-\pi)} \equiv \widehat{\pi}_l^R < \pi$.

We must thus consider six different cases:

1. $\widehat{\pi} = 1$ and $\widehat{\lambda} = 0$;

2. $\widehat{\pi} = 1$ and $\widehat{\lambda} = 1$;

3. $\widehat{\pi} = 0$;

4. $\widehat{\pi} = \widehat{\pi}_l^L$ and $\widehat{\lambda} = 0$;

5. $\widehat{\pi} = \widehat{\pi}_l^R$ and $\widehat{\lambda} = 1$;

6. $\widehat{\lambda} = \lambda$ and $\widehat{\pi} = \pi$.

We will assume that the shock realization favors the right-wing party (a similar analysis establishes the result for the case in which the shock favors L).

Case 1: $\widehat{\pi} = 1$ and $\widehat{\lambda} = 0$.

Here, the voter learns that the policy dimension is relevant for her welfare and her optimal policy is 1, which is also R's bliss point. Party R is favored by the shock, by assumption, therefore if R proposes policy 1 it wins with probability 1. Thus, it must be the case that in equilibrium $x = 1$. To establish a contradiction, conjecture an equilibrium in which $x \neq 1$. Then, it must be the case that R is proposing a policy $x \neq 1$. However, R can always move to policy 1 and win for sure, which is a profitable deviation.

Case 2: $\widehat{\pi} = 1$ and $\widehat{\lambda} = 1$.

The voter learns that the policy dimension is relevant for her welfare, and her optimal policy is -1. Denote x' the policy that makes the voter indifferent between electing the valence-favored party R and getting policy x', and electing L and getting policy -1. x' thus solves $0 = \delta - (-1-x)^2$, which yields $x' = -1 + \sqrt{\delta}$.

Straightforwardly, the highest utility L can offer to the voter is by proposing -1. Therefore, by definition of x', R can always win by proposing $x = \min\{1, x'\}$, regardless of what policy L proposes. If $x' > 1$, in equilibrium R must win with probability 1 by proposing policy 1. For any other possible equilibrium, R can deviate closer to 1 without decreasing its probability of winning. If $x' < 1$, then in equilibrium L must propose -1 and R must propose x', and the voter must break indifference by electing the valence-favored R. Notice that the voter must use this indifference breaking rule in equilibrium, as otherwise party R always has a profitable deviation to make an arbitrarily small move to the left.

Case 3: $\widehat{\pi} = 0$.

The voter is indifferent between all policies, and therefore votes based solely on the valence dimension. R is favored by the valence realization, by assumption, therefore in equilibrium always proposes its preferred policy and wins with probability 1.

Case 4: $\widehat{\pi} = \widehat{\pi}_l^L$ and $\widehat{\lambda} = 0$.

This case is similar to case 1. Even though the voter does not learn whether the policy dimension is relevant, she learns that, if it is, her optimal policy must take value 1. Because R is favored by the valence shock, it can always win by proposing its preferred policy 1, and no other outcome can be sustained in equilibrium.

Case 5: $\widehat{\pi} = \widehat{\pi}_l^R$ and $\widehat{\lambda} = 1$.

Symmetric to the previous case, the voter does not learn whether the policy dimension is relevant, but she learns that, if it is, her optimal policy must take value -1. This case is analogous to case 3, and the proof proceeds in the same way, with x' now solving $0 = \delta - \widehat{\pi}_l^R(-1-x)^2$.

Case 6: $\widehat{\pi} = \pi$ and $\widehat{\lambda} = \lambda$.

This case is similar to cases 2 and 5, but x' now solves $-\pi\Big(\lambda(-1-x^L)^2 + (1-\lambda)(1-x^L)^2\Big) = \delta - \pi\Big(\lambda(-1-x)^2 + (1-\lambda)(1-x)^2\Big)$, with $x^L = 1 - 2\lambda$. □

Claim 2. *There exists a unique* $\widetilde{\delta}$ *s.t. an equilibrium in which some party makes a supporting argument can be sustained only if* $\delta > \widetilde{\delta}$.

Proof. Focusing w.l.o.g. on party L, there are three argument profiles we must consider:

1. $(a^R = s, a^L = s)$;
2. $(a^R = v, a^L = s)$;
3. $(a^R = r, a^L = s)$.

We will establish, for each profile, that a necessary condition for L to have no profitable deviation is that δ is sufficiently high. To reduce the number of cases under consideration, we will assume that $\delta < \max\{4\widehat{\pi}_l^L, 4\widehat{\pi}_l^R\}$. Further, recall that $\nu_R \in (0,1)$ is the probability that the valence shock favors the right-wing party.

Case 1: $(a^R = s, a^L = s)$**.**

Recall that the shock will favor one party or the other, so that the probability of it favoring L is simply $1 - \nu_R$. Then, given the policies implemented in equilibrium, as characterized in Claim 1, in the conjectured equilibrium, L gets expected payoff

$$-(1-\pi)\nu_R 4 - \pi\lambda\nu_R\delta - \pi(1-\lambda)\nu_R 4 - \pi(1-\lambda)(1-\nu_R)(2-\sqrt{\delta})^2.$$

Consider now a deviation. Recall that, in our setting, a supporting argument from one party on dimension j provides the same information as a refuting argument from the other party on the same dimension. Then, given $a^R = s$, for party L a deviation to a vacuous or refuting argument is payoff-equivalent, and, assuming $\delta < 4\widehat{\pi}_l^R$, yields

$$-\pi(1-\lambda)\nu_R 4 - \pi(1-\lambda)(1-\nu_R)(2-\sqrt{\delta})^2 - (1-\pi(1-\lambda))\nu_R \frac{\delta}{\widehat{\pi}_l^R}.$$

Thus, the deviation is profitable iff

$$\begin{aligned}&-(1-\pi)\nu_R 4-\pi\lambda\nu_R\delta-\pi(1-\lambda)\nu_R 4-\pi(1-\lambda)(1-\nu_R)(2-\sqrt{\delta})^2 \\ &+\pi(1-\lambda)\nu_R 4+\pi(1-\lambda)(1-\nu_R)(2-\sqrt{\delta})^2+(1-\pi(1-\lambda))\nu_R\frac{\delta}{\widehat{\pi}_l^R}<0.\end{aligned} \tag{12}$$

the LHS is continuous and strictly increasing in δ, and the condition is satisfied when $\delta = 0$ and fails at $\delta = 4\widehat{\pi}_l^R$.

Suppose instead that $\delta \in (4\widehat{\pi}_l^R, 4\widehat{\pi}_l^L)$. Then, for party L a deviation to a vacuous or refuting argument yields expected payoff

$$-\pi(1-\lambda)\nu_R 4-\pi(1-\lambda)(1-\nu_R)(2-\sqrt{\delta})^2-(1-\pi(1-\lambda))\nu_R 4,$$

and is profitable iff

$$\begin{aligned}&-(1-\pi)\nu_R 4-\pi\lambda\nu_R\delta-\pi(1-\lambda)\nu_R 4-\pi(1-\lambda)(1-\nu_R)(2-\sqrt{\delta})^2 \\ &+\pi(1-\lambda)\nu_R 4+\pi(1-\lambda)(1-\nu_R)(2-\sqrt{\delta})^2+(1-\pi(1-\lambda))\nu_R 4<0.\end{aligned} \tag{13}$$

This condition is never satisfied. Thus, there must exist a unique threshold strictly smaller than $4\widehat{\pi}_l^R$ s.t. L has no profitable deviation iff δ above the threshold.

Case 2: $(a^R = v, a^L = s)$**.** Assuming δ is sufficiently small to guarantee that the equilibrium platforms are always interior, in the conjectured equilibrium L's expected payoff is

$$-\pi\lambda\nu_R\delta-(1-\pi\lambda)\nu_R 4-(1-\pi\lambda)(1-\nu_R)(2-\sqrt{\frac{\delta}{\widehat{\pi}_l^L}})^2.$$

Recall that in this claim we want to establish necessary conditions to sustain an equilibrium with supporting arguments. Thus, it is enough to show that, for δ sufficiently low, one of the players has a profitable deviation. For L, a deviation to a vacuous argument yields

$$-\nu_R(1-2\lambda+\sqrt{\frac{\delta}{\pi}}+1)^2-(1-\nu_R)(1-2\lambda-\sqrt{\frac{\delta}{\pi}}+1)^2,$$

and is therefore profitable iff

$$\begin{aligned}&-\nu_R(1-2\lambda+\sqrt{\frac{\delta}{\pi}}+1)^2-(1-\nu_R)(1-2\lambda-\sqrt{\frac{\delta}{\pi}}+1)^2\\&+\pi\lambda\nu_R\delta+(1-\pi\lambda)\nu_R4+(1-\pi\lambda)(1-\nu_R)(2-\sqrt{\frac{\delta}{\widehat{\pi}_l^L}})^2>0.\end{aligned} \tag{14}$$

Observe that the LHS of this inequality is continuous in δ, and the condition is satisfied at $\delta=0$. Thus, there exists a threshold of δ s.t. L always has a profitable deviation if δ is below this threshold.

Case 3: $(a^R=r, a^L=s)$**.** Assuming $\delta<4\widehat{\pi}_l^L$, in the conjectured equilibrium L's expected payoff is

$$-\pi\lambda\nu_R\delta-(1-\pi\lambda)\nu_R4-(1-\pi\lambda)(1-\nu_R)(2-\sqrt{\frac{\delta}{\widehat{\pi}_l^L}})^2.$$

A deviation to a refuting argument yields

$$-(1-\pi)\nu_R4-\pi\lambda\nu_R\delta-\pi(1-\lambda)\nu_R4-\pi(1-\lambda)(1-\nu_R)(2-\sqrt{\delta})^2,$$

and is profitable unless

$$\begin{aligned}&-\pi\lambda\nu_R\delta-(1-\pi\lambda)\nu_R4-(1-\pi\lambda)(1-\nu_R)(2-\sqrt{\frac{\delta}{\widehat{\pi}_l^L}})^2\\&(1-\pi)\nu_R4+\pi\lambda\nu_R\delta+\pi(1-\lambda)\nu_R4+\pi(1-\lambda)(1-\nu_R)(2-\sqrt{\delta})^2>0.\end{aligned}$$

The LHS of this inequality is concave in δ on $[0,4\widehat{\pi}_l^L]$, and is always satisfied at $\delta=4\widehat{\pi}_l^L$ and never satisfied at $\delta=0$. Thus, as above, there must exist a unique threshold of δ on $[0,4\widehat{\pi}_l^L]$ s.t. the deviation is profitable iff δ is below that threshold.

Therefore, in each of the three cases, there exists a threshold in δ s.t. L always has profitable deviation if δ is smaller than the corresponding threshold. Following the same logic as above one can establish the same result for party R. The threshold $\widetilde{\delta}$ is then characterized by identifying the most binding no-deviation condition (across all three cases and both players), and choosing the δ that satisfies the condition with equality. □

Claim 3. *There exists a unique pair* $(\widehat{\delta}, \pi^{\dagger})$ *s.t., if* $\delta > \widehat{\delta}$ *and* $\pi < \pi^{\dagger}$*, then, in the unique equilibrium both parties present supporting arguments.*

Proof. First, we know from the above analysis that there exists a unique threshold in δ s.t. an equilibrium in which both parties present supporting arguments exists iff δ is above this cutoff (see conditions (12) and (13)). This also implies that above this cutoff we cannot sustain equilibria in which only one party presents a supporting argument, as the other party will have a profitable deviation to a supporting argument.

Next, we show that an equilibrium with two vacuous arguments cannot be sustained when δ is sufficiently high, as each party has a profitable deviation to a supporting one. To see this, suppose first that $\lambda > \frac{1}{2}$. This implies that $\widehat{\pi}_l^L < \widehat{\pi}_l^R$. Furthermore, $4\pi\lambda^2 < \widehat{\pi}_l^R$. Recall that we are assuming $\delta < \max\{4\widehat{\pi}_l^L, 4\widehat{\pi}_l^R\}$. Suppose then that $\delta \in (max\{4\widehat{\pi}_l^L, 4\pi\lambda^2\}, 4\widehat{\pi}_l^R)$. Then, using the equilibrium policies characterized in Claim 1, L's no-deviation condition (condition (14))[2] becomes

$$-\nu_R 4 + \pi\lambda\nu_R\delta + (1-\pi\lambda)\nu_R 4 > 0,$$

which is never satisfied. Thus, the conjectured equilibrium can never be sustained. A similar analysis establishes that R has a profitable deviation from the conjectured equilibrium when $\delta \in (max\{4\widehat{\pi}_l^R, 4\pi(1-\lambda)^2\}, 4\widehat{\pi}^L)$ and $\lambda < \frac{1}{2}$.

[2]Condition (14) assumes that δ is sufficiently small that the equilibrium platforms are always interior. Here, we instead assume that it is sufficiently high that $1-2\lambda-\sqrt{\frac{\delta}{\pi}} < -1$, $1-2\lambda+\sqrt{\frac{\delta}{\pi}} > 1$ and $1-\sqrt{\frac{\delta}{\widehat{\pi}_l^L}} < -1$.

Next, conjecture an equilibrium where both parties present a refuting argument. If $\delta < 4\widehat{\pi}_l^L$, the no-deviation condition for L is

$$-(1-\pi)\nu_R 4 - \pi\lambda\nu_R\delta - \pi(1-\lambda)\nu_R 4 - \pi(1-\lambda)(1-\nu_R)(2-\sqrt{\delta})^2$$
$$+\pi\lambda\nu_R\delta + (1-\pi\lambda)\nu_R 4 + (1-\pi\lambda)(1-\nu_R)(2-\sqrt{\frac{\delta}{\widehat{\pi}_l^L}})^2 \geq 0.$$

The LHS of the above condition is continuous in δ, and the condition is never satisfied at $\delta = 4\widehat{\pi}_l^L$. If instead $\delta \in (4\widehat{\pi}_l^L, 4\widehat{\pi}_l^R)$, then the no-deviation condition for L is

$$-(1-\pi)\nu_R 4 - \pi\lambda\nu_R\delta - \pi(1-\lambda)\nu_R 4 - \pi(1-\lambda)(1-\nu_R)(2-\sqrt{\delta})^2 + \pi\lambda\nu_R\delta \geq 0,$$

and is never satisfied. Therefore, there must exist a threshold strictly smaller than $4\widehat{\pi}_l^L$ s.t. the condition always fails above this threshold and the equilibrium cannot be sustained.

Finally, conjecture an equilibrium where L presents a vacuous argument and R a refuting one. Assume first $\lambda < \frac{1}{2}$, which implies $4\widehat{\pi}_l^L > 4\widehat{\pi}_l^R$ and therefore $\delta < 4\widehat{\pi}_l^L$. Then, a deviation to a refuting argument is profitable for L whenever:

$$-\pi\lambda\nu_R\delta - (1-\pi\lambda)\nu_R 4 - (1-\pi\lambda)(1-\nu_R)(2-\sqrt{\frac{\delta}{\widehat{\pi}_l^L}})^2 <$$
$$-(1-\pi)\nu_R 4 - \pi\lambda\nu_R\delta - \pi(1-\lambda)\nu_R 4 - \pi(1-\lambda)(1-\nu_R)(2-\sqrt{\delta})^2.$$

This condition reduces to

$$-(1-\pi\lambda)(2-\sqrt{\frac{\delta}{\widehat{\pi}_l^L}})^2 < -\pi(1-\lambda)(2-\sqrt{\delta})^2,$$

which always holds for a sufficiently small π.

Assume instead $\lambda > \frac{1}{2}$, which implies $4\widehat{\pi}_l^L < 4\widehat{\pi}_l^R$ and therefore $\delta < 4\widehat{\pi}_l^R$. Suppose $\delta \in (4\widehat{\pi}_l^L, 4\widehat{\pi}_l^R)$.

A deviation to a supporting argument is profitable for R iff

$$\begin{aligned} & -\nu_R \pi\lambda(-2+\sqrt{\delta})^2 - (1-\nu_R)\pi\lambda 4 - (1-\nu_R)(1-\pi\lambda)4 \\ < & -\nu_R(1-\pi(1-\lambda))(2-\sqrt{\frac{\delta}{\widehat{\pi}_l^R}})^2 - (1-\nu_R)\pi(1-\lambda)\delta - (1-\nu_R)(1-\pi(1-\lambda))4, \end{aligned}$$

which is always satisfied at $\delta = 4\widehat{\pi}_l^R$. By continuity in this range, there is a cutoff strictly smaller than $4\widehat{\pi}_l^R$ s.t. the deviation is profitable if δ is above this cutoff.

A similar analysis establishes the results for a conjecture in which R presents a vacuous argument and L a refuting one.

Thus, there must exist cutoffs $\widehat{\delta} \geq \widetilde{\delta}$ s.t. and $\pi^\dagger$ s.t. when $\delta > \widehat{\delta}$ and $\pi < \pi^\dagger$, the game has a unique equilibrium, in which both parties present supporting arguments. □

Proposition 2C. *There exist unique* $\underline{\delta} \geq \tilde{\delta}$ *and* $\overline{\delta} < \widehat{\delta}$ *s.t. an equilibrium in which one party presents a refuting argument and the other a supporting argument exists if and only if* $\delta \in [\underline{\delta}, \overline{\delta}]$.

Proof. Conjecture an equilibrium where R presents a refuting argument and L a supporting one. From the proof of Proposition 4, L has no profitable deviation if and only if

$$-(1-\pi\lambda)(2-\sqrt{\frac{\delta}{\widehat{\pi}_l^L}})^2 + \pi(1-\lambda)(2-\sqrt{\delta})^2 > 0.$$

and there must exist a cutoff $\underline{\delta}$ s.t. the condition is satisfied if δ is above $\underline{\delta}$ and fails otherwise.

Next, consider party R. Applying Claim 1, R has no profitable deviation if and only if

$$\begin{aligned} & -\nu_r\pi\lambda(2-\sqrt{\delta})^2 - (1-\nu_r)\pi\lambda 4 - (1-\nu_r)(1-\pi\lambda)\frac{\delta}{\widehat{\pi}_l^L} > \\ & -\nu_r\pi\lambda(2-\sqrt{\delta})^2 - (1-\nu_r)(1-\pi(1-\lambda))4 - (1-\nu_r)\pi(1-\lambda)(2-\sqrt{\delta})^2, \end{aligned}$$

which reduces to

$$-(1-\pi\lambda)\frac{\delta}{\widehat{\pi}_l^L} + (1-\pi)4 + \pi(1-\lambda)(2-\sqrt{\delta})^2 > 0.$$

The RHS is decreasing in δ, always satisfied at $\delta = 0$ and never satisfied at $\delta = 4\widehat{\pi}_l^L$. There must exist a cutoff $\overline{\delta}$ s.t. the condition is satisfied if δ is below $\overline{\delta}$ and fails otherwise.

To conclude the proof, we must show that there exist conditions under which $\overline{\delta} > \underline{\delta}$. It can be easily verified that this holds when $\lambda = \pi = \frac{1}{2}$. Therefore, it must also hold for a range of values around this point.

Similar steps establish the conditions for an equilibrium where L presents a refuting argument and R a supporting one. □

In the following analysis, we consider the interpretation in which δ captures the voter's (realized) preferences over a second policy dimension, on which that the parties have preferences but no flexibility over. This is technically equivalent to assuming that, in addition to the policy payoff over the main dimension, parties pay a cost K from losing the election (i.e., a cost equivalent to the difference in their utility from the implementation of their policy and their opponent's on the inflexible dimension).

Proposition 3C. *Suppose that, in addition to their policy payoffs, each party $j \in \{L, R\}$ pays a cost K from losing the election. The equilibrium of the argumentation stage is as characterized in Proposition 4.*

Proof. In order to prove this result, we must establish the following claim:

Claim 4.

1. *The equilibrium policy is as characterized in Claim 1 in the proof of Proposition 4;*

2. *the equilibrium probability of each party winning the election is independent of the argument profile presented in the argumentation stage.*

Proof. We proceed by cases as in the proof of Claim 1 in Proposition 4. Suppose that the realization of δ favors party R. Symmetric results hold for the case in which the shock favors L.

Case 1: $\widehat{\pi} = 1$ and $\widehat{\lambda} = 0$.

By the logic of the proof of Claim 1, if R proposes policy 1 it wins with probability 1. Thus, it must be the case that, in equilibrium, $x = 1$ and R wins with certainty.

Case 2: $\widehat{\pi} = 1$ **and** $\widehat{\lambda} = 1$**.**

By the logic of the proof of Claim 1, R always wins in equilibrium, and proposes $x = \min\{1, x'\}$, where $x' = -1 + \sqrt{\delta}$.

Case 3: $\widehat{\pi} = 0$**.**

The voter is indifferent between all policies, and therefore votes based solely on the valence dimension. R is favored by the valence realization, by assumption, therefore in equilibrium always proposes its preferred policy and wins with probability 1.

Case 4: $\widehat{\pi} = \widehat{\pi}_l^L$ **and** $\widehat{\lambda} = 0$**.**

Because R is favored by the valence shock, it wins with certainty by proposing its preferred policy 1.

Case 5: $\widehat{\pi} = \widehat{\pi}_l^R$ **and** $\widehat{\lambda} = 1$**.**

This case is analogous to case 2, and the proof proceeds in the same way, with x' now solving $0 = \delta - \widehat{\pi}_l^R(-1 - x)^2$.

Case 6: $\widehat{\pi} = \pi$ **and** $\widehat{\lambda} = \lambda$**.**

This case is analogous to cases 2 and 5, but x' now solves $-\pi\Big(\lambda(-1-x^L)^2+(1-\lambda)(1-x^L)^2\Big) = \delta - \pi\Big(\lambda(-1-x)^2 + (1-\lambda)(1-x)^2\Big)$, with $x^L = 1 - 2\lambda$.

Thus, if the realization of δ favors R, the equilibrium policy is $\min\{1, 1 - 2\widehat{\lambda} + \sqrt{\frac{\delta}{\widehat{\pi}}}\}$, as characterized in Claim 1, and R wins with certainty. Similarly, we can show that if the realization of δ favors L, the equilibrium policy is $\max\{-1, -1 + 2\widehat{\lambda} + \sqrt{\frac{\delta}{\widehat{\pi}}}\}$, and L wins with certainty. □

Claim 4 implies that the expected continuation value from any argument profile is identical in the two versions of the model with frictions, therefore the equilibrium of the argumentation stage must be identical as well. □

D Proofs for Extension with Correlated Types

In this section, we introduce some correlation in the voter's types across dimensions. Specifically, we assume that the players share common prior beliefs that with probability $p \in (0, 1)$ the voter's type is drawn from a joint distribution with correlation 1 (i.e., $\theta_A = \theta_B$) and parameters π and λ. With the complement probability $1 - p$ the types on the two dimensions are uncorrelated, and drawn from distributions with parameters λ_j and π_j as in the baseline. To reduce notational burden but without much loss of substance, we further set $\lambda_A = \lambda_B = \lambda$ and $\pi_A = \pi_B := \pi$.

The assumption that the dimensions are potentially correlated has a crucial consequence: if the voter receives a non-vacuous argument on dimension j, the resonance event of this argument will inform the voter's beliefs over *both* dimensions j and $-j$. Nonetheless, the insights of our baseline model remain robust.

Lemma 1D. *If a non-vacuous argument is presented on dimension j, then, in equilibrium,*

1. *policy implemented on j does not depend on arguments presented on dimension $k \neq j$;*
2. *arguments presented on dimension $k \neq j$ satisfy either $a_k^L = a_k^R = v$ or $a_k^L = a_k^R = r$;*
3. *a supporting argument is presented on j only if $a_k^L = a_k^R = v$.*

Proof. To establish (1), first notice that the equilibrium of the platform game remains unchanged from the baseline:

- If $\widehat{\pi}_j > 0$, then implemented policy on dimension j is equal to the policy that maximizes the voter's expected utility, given her posterior beliefs;
- If $\widehat{\pi}_j = 0$, then the implemented policy on dimension j is equal to $\tilde{x}_j^R$ with probability $\frac{1}{2}$, and equal to $\tilde{x}_j^L$ with probability $\frac{1}{2}$.

Next, notice that, since $p \in (0, 1)$, the resonance event of arguments on dimension j cannot induce the voter to believe that $\theta_{-j} = \emptyset$ with probability 1.

Further, recall from our discussion in the main model that, whenever at least a non-vacuous argument is presented on dimension j, the voter always reaches degenerate posteriors (either $\widehat{\pi}_j = 0$, or $\widehat{\lambda}_j \in \{0, 1\}$).

Suppose two refuting arguments are presented on dimension j. If neither resonates, $\widehat{\pi}_j = 0$, and the voter is indifferent across all policies. If one resonates and the other does not, $\widehat{\pi}_j > 0$ and $\widehat{\lambda}_j$ is either 0 or 1, depending on the argument that resonates. Similar results apply if two supporting arguments are presented.

Suppose instead a single non-vacuous argument is presented on dimension j. Again, the voter's posterior $\widehat{\lambda}_j$ is always degenerate, at 0 or 1. If R presents a refuting argument and it resonates, $\widehat{\lambda}_j = 0$. If it does not resonate, $\widehat{\lambda}_j = 1$. If R presents a supporting argument and it resonates, $\widehat{\lambda}_j = 0$. If it does not resonate, $\widehat{\lambda}_j = 1$. Symmetric results apply to party L.

Thus, whenever the voter receives at least one non-vacuous argument on dimension j, either she learns all policies on j are equivalent, $\widehat{\pi}_j = 0$, or she updates that $\widehat{\pi}_j > 0$ and her posterior $\widehat{\lambda}_j$ is unaffected by the arguments presented on $-j$. This concludes the proof.

Next, from (1), and the results of the baseline model, we have that (2) must be satisfied.

Finally, from (2), if $s \in \{a_k^L, a_k^R\}$ in equilibrium, then $a_j^L = a_j^R = v$, establishing (3). □

Next, we must characterize how the voter's preferences on dimension j depend on her posteriors on dimension $-j$, in absence of non-vacuous j-dimension arguments. Recall that x_j^v is the policy that maximizes the voter's expected utility on dimension j, given her posterior beliefs reached as a result of the argumentation stage. Then, we have

Lemma 2D. *Suppose that $a_j^L = a_j^R = v$. The voter's induced preferences on dimension j, x_j^v, have the following properties*

- *If the voter learns that $\theta_{-j} = 1$, then* $x_j^v(\theta_{-j} = 1) = \frac{p+(1-p)\pi(1-2\lambda)}{p+(1-p)\pi}$
- *If the voter learns that $\theta_{-j} = -1$, then* $x_j^v(\theta_{-j} = -1) = \frac{-p+(1-p)\pi(1-2\lambda)}{p+(1-p)\pi}$
- *If the voter learns that $\theta_{-j} = \emptyset$, then* $x_j^v(\theta_{-j} = \emptyset) = 1 - 2\lambda$

- *If the voter learns that* $\theta_{-j} \neq -1$*, then* $x_j^v(\theta_{-j} \neq -1) = \frac{p\tilde{\pi}(\theta_{-j}\neq-1)+(1-p)\pi(1-2\lambda)}{p\tilde{\pi}(\theta_{-j}\neq-1)+(1-p)\pi}$*, where where* $\tilde{\pi}(\theta_{-j} \neq -1) \leq \pi$ *is the posterior probability that* $\theta_j \neq \emptyset$ *conditional on the dimensions being correlated*

- *If the voter learns that* $\theta_{-j} \neq 1$*, then* $x_j^v(\theta_{-j} \neq 1) = \frac{-p\tilde{\pi}(\theta_{-j}\neq1)+(1-p)\pi(1-2\lambda)}{p\tilde{\pi}(\theta_{-j}\neq1)+(1-p)\pi}$*, where* $\tilde{\pi}(\theta_{-j} \neq 1) \leq \pi$ *is the posterior probability that* $\theta_j \neq \emptyset$ *conditional on the dimensions being correlated*

Proof. Follows from applying Bayes rule and solving the voter's maximization problem. □

Proposition 4D. *There is no equilibrium in which a party presents a supporting argument.*

Proof. From Lemma 1D part (3), if there exists an equilibrium with a supporting argument, then there must be a dimension on which there are only vacuous arguments in that equilibrium. Without loss of generality, call that dimension 2, suppose $(a_2^L, a_2^R) = (v, v)$, and consider the possibilities $a_1 \in \{(s, v), (v, s), (s, s), (r, s), (s, r)\}$ to complete an equilibrium strategy profile.

1. Consider $a_1 = (s, v)$. Applying Lemma 2D,

$$\begin{aligned} E[U^L((s,v),a_2] = &- 4(1-\pi\lambda) - 4\pi\lambda[\frac{(1-p)\pi(1-\lambda)}{p+(1-p)\pi}]^2 \\ &- 4(1-\pi\lambda)[\frac{p(1-\lambda)+(1-p)(1-\lambda)(1-\pi\lambda)}{p(1-\lambda)+(1-p)(1-\pi\lambda)}]^2, \end{aligned} \tag{15}$$

   and

$$\begin{aligned} E[U^L((r,v),a_2] = &- 4\pi(1-\lambda) - 4\pi(1-\lambda)[\frac{p+(1-p)\pi(1-\lambda)}{p+(1-p)\pi}]^2 \\ &- 4(1-\pi(1-\lambda))[\frac{(1-p)(1-\lambda)(1-\pi(1-\lambda))}{p\lambda+(1-p)(1-\pi(1-\lambda))]}]^2. \end{aligned} \tag{16}$$

   Because $E[U^L((r, v), a_2)] > E[U^L((s, v), a_2)] \forall p \in (0, 1) \forall\lambda \in (0, 1) \forall\pi \in (0, 1)$, $((s, v),(v, v))$ is never an equilibrium. A symmetric argument for $a_1 = (v, s)$[3] showing that party R prefers to deviate to r establishes that $((v, s),(v, v))$ is never an equilibrium.

[3]We are using the notation a_j to denote the profile of arguments presented by the parties on dimension j.

2. Consider $a_1 = (s, s)$. Applying Lemma 2D,

$$E[U^L((s,s),a_2)] = -4\pi(1-\lambda) - 2(1-\pi) - 4(1-\pi)(1-\lambda)^2 - 4\pi\lambda[\frac{(1-p)\pi(1-\lambda)}{p+(1-p)\pi}]^2 \quad (17)$$

$$-4\pi(1-\lambda)[\frac{p+(1-p)\pi(1-\lambda)}{p+(1-p)\pi}]^2. \quad (18)$$

Note that $E[U^L((r,s),a_2)] = E[U^L((v,s),a_2)] = E[U^L((r,v),a_2)]$, which is characterized in (16). Because $E[U^L((r,s),a_2)] > E[U^L((s,s),a_2)] \forall p \in (0,1) \forall \lambda \in (0,1) \forall \pi \in (0,1)$, $((s,s),(v,v))$ is never an equilibrium.

3. Consider $a_1 = (s, r)$. Notice that $E[U^L((r,r),a_2)] = E[U^L((s,s),a_2)]$, which is characterized in (17), and $E[U^L((s,r),a_2)] = E[U^L((s,v),a_2)]$, which is characterized in (15). Because $E[U^L((r,r),a_2)] > E[U^L((s,r),a_2)] \forall p \in (0,1) \forall \lambda \in (0,1) \forall \pi \in (0,1)$, $((s,r),(v,v))$ is never an equilibrium. A symmetric argument for $a_1 = (r, s)$, showing that party R prefers to deviate to r establishes that $((r,s),(v,v))$ is never an equilibrium. □

E Proofs for Richer Argument Space

In this section, we enrich the parties' argument space. In particular, we assume that parties can make a *salience argument*, which aims to persuade the voter that she should or should not care about a specific dimension; a refuting, supporting, or vacuous argument, as in the baseline model; or a combination of multiple arguments.[4]

We assume that each party pays an arbitrarily small cost for each non-vacuous argument it presents.[5]

Proposition 5E. *There exist unique thresholds $\widetilde{\pi}_j(\lambda_j)$, $\underline{\lambda}_j$, and $\overline{\lambda}_j$, $\underline{\lambda}_j < \frac{1}{2} < \overline{\lambda}_j$, such that*

- *if $\pi_j > \widetilde{\pi}_j(\lambda_j)$, then, in any Pareto-undominated equilibrium, on dimension j*

[4]To use our cups analogy, we allow each party to flip any subset of cups.

[5]This has no effect on the results of the baseline model, but ensures equilibrium uniqueness.

- *both parties present vacuous arguments if* $\lambda_j \in (\underline{\lambda}_j, \overline{\lambda}_j)$;
- *L presents a salience argument and R a vacuous one if* $\lambda_j \leq \underline{\lambda}_j$;
- *R presents a salience argument and L a vacuous one if* $\lambda_j \geq \overline{\lambda}_j$;

- *if* $\pi_j < \widetilde{\pi}_j(\lambda_j)$, *then in any equilibrium both parties present refuting arguments.*

Proof. Recall that any action profile where parties present two or more arguments allows for full learning for the voter. This also holds if the arguments are presented by the same party. Further, notice that our assumption on the arbitrarily small cost of presenting non-vacuous arguments implies that informationally redundant arguments cannot be sustained in equilibrium. These observations, combined with the results of the baseline model, leave six equilibrium candidates, which we consider below.

1. Party i presents a fully informative argument, and party $-i$ *presents a vacuous argument.*

This equilibrium cannot be sustained, as a deviation to a refuting argument is always profitable. To establish a contradiction, conjecture an equilibrium in which R presents a fully informative argument and L presents a vacuous one (an analogous argument applies to the symmetric conjecture). R's expected payoff in the conjectured equilibrium is $-\pi_j\lambda_j 4 - (1-\pi_j)4\frac{1}{2}$. A deviation to $a_j^R = r$ yields a strictly higher expected payoff $-\pi_j\lambda_j 4$.

2. Party i presents a salience argument and party $-i$ *a supporting one.*

The same logic as in the previous case implies that such an equilibrium cannot be sustained, as $-i$ has profitable deviation to presenting a vacuous argument. Suppose that L presents a supporting argument and R a salience one. In the conjectured equilibrium, R's expected payoff is $-\pi_j\lambda_j 4 - (1-\pi_j)4\frac{1}{2}$. A deviation to a vacuous argument yields a strictly higher expected payoff $-\pi_j\lambda_j 4$. Similar analysis establishes the result for the symmetric case.

3. Party i presents a salience argument and party $-i$ *a refuting one.*

Suppose R presents a salience argument and L a refuting one (the analysis and the conclusion in the symmetric case are analogous). From the analysis of the baseline model we know that R has no profitable deviation from the conjectured strategy (see (7)). Party L's expected payoff in the conjectured equilibrium is

$$-\pi_j(1-\lambda_j)4-(1-\pi_j)\frac{1}{2}4.$$

Given R's salience argument, a deviation to L still allows for full learning, and therefore is payoff-irrelevant. Suppose instead L deviates to a vacuous or salience argument. The deviation yields expected payoff

$$-\pi_j\big(-1-(1-2\lambda_j)\big)^2-(1-\pi_j)\frac{1}{2}4,$$

which reduces to

$$-\pi_j 4(1-\lambda_j)^2-(1-\pi_j)\frac{1}{2}4,$$

and is therefore always profitable.

4. Party i presents a salience argument and party $-i$ a vacuous one.

Consider a conjecture in which L presents a salience argument and R a vacuous one. From the previous analysis of equilibrium candidates 2 and 3, R has no profitable deviation. In the conjectured equilibrium, L's expected payoff is

$$-\pi_j 4(1-\lambda_j)^2-(1-\pi_j)\frac{1}{2}4.$$

We know from the above analysis that a deviation to a fully informative argument is not profitable. Suppose instead L deviates to a refuting argument. This yields expected utility $-\pi_j(1-\lambda_j)4$. Thus, the deviation is profitable iff

$$-\pi_j 4(1-\lambda_j)^2-(1-\pi_j)\frac{1}{2}4<-\pi_j(1-\lambda_j)4. \tag{19}$$

This establishes a cutoff $\tilde{\pi}_j$ s.t. the deviation is profitable iff π_j is below this cutoff.

A deviation to a supporting argument yields L expected payoff $-(1-\pi_j\lambda_j)4$ and is therefore never profitable.

Finally, a deviation to a vacuous argument yields party L expected payoff $-4(1-\lambda_j)^2$ and is profitable iff

$$-\pi_j 4(1-\lambda_j)^2 - (1-\pi_j)\frac{1}{2}4 < -4(1-\lambda_j)^2.$$

Rearranging and canceling terms, we obtain

$$(1-\lambda_j)^2 < \frac{1}{2}. \tag{20}$$

This establishes a cutoff $\underline{\lambda}_j$ s.t. the deviation is profitable iff λ_j is above this cutoff.

From (19) and (20), we have that there exist unique $\tilde{\pi}_j$ and $\underline{\lambda}_j < \frac{1}{2}$ s.t. an equilibrium in which L presents a salience argument and R presents a vacuous one exists iff $\pi_j > \tilde{\pi}_j^L$ and $\lambda_j < \underline{\lambda}_j$.

A similar argument establishes the result for an equilibrium in which R presents a salience argument and L a vacuous one: there exists a unique $\overline{\lambda}_j$ and $\tilde{\pi}_j^R$ s.t. the conjectured equilibrium exists iff $\pi_j > \tilde{\pi}_j^R$ and $\lambda > \overline{\lambda}_j > \frac{1}{2}$.

5. Both parties present refuting arguments.

A deviation from the conjectured equilibrium to salience or complex arguments is still fully revealing and thus payoff-irrelevant, therefore the availability of these arguments has no effect on the equilibrium analysis. As a consequence, our results from the baseline continue to apply: this equilibrium can always be sustained, but it does not survive our selection criterion when π_j is sufficiently high.

6. Both parties present vacuous arguments.

Following the previous analysis of equilibrium candidate 4, a unilateral deviation to a salience argument is profitable for L iff $\lambda_j < \underline{\lambda}_j$. Similarly, a unilateral deviation to a salience argument is profitable for R when λ_j is sufficiently large, $\lambda_j > \overline{\lambda}_j$. Finally, the baseline model establishes that a necessary condition to sustain this equilibrium is that π_j is sufficiently large (see conditions (8) and (9)).

To conclude the proof, we must only establish that, when an equilibrium with a salience argument exists, it yields higher expected payoff for both parties than does the equilibrium in which both parties present a refuting argument. Consider party L. An equilibrium in which only a salience argument is presented yields

$$-\pi_j 4(1-\lambda_j)^2 - (1-\pi_j)\frac{1}{2}4.$$

An equilibrium in which both present refuting arguments yields for the L a strictly lower expected utility

$$-\pi_j 4(1-\lambda_j) - (1-\pi_j)\frac{1}{2}4.$$

The corresponding logic shows that R's expected utility is higher under a salience argument. □

F Proof for Extension with Cost of Inconsistency

In this section, we analyze an extension incorporating the intuition that voters may punish political parties for inconsistencies between the arguments they present and the platforms they commit to. Specifically, we require that a party that presents a non-vacuous argument must then adopt the platform aligned with the argument in the next stage. Thus, if the left-wing (right-wing) party advances a non-vacuous argument to persuade the voter to support platform -1 $(+1)$, the party is then forced to commit to platform -1 $(+1)$ in the platform game.[6] This, in effect, captures a stark version of the voter's preference for consistency, where the implicit punishment for inconsistency is prohibitive for political parties.

Our analysis shows that, while the consistency constraint alters the parties' equilibrium platforms, each argument profile nonetheless results in the same lottery of policies as it does in the baseline model.

Lemma 1F. *In any equilibrium, the parties' platform choices satisfy*

[6] Notice that refuting and supporting arguments equally constrain political parties. It could be argued, however, that refuting arguments might preserve some flexibility in platform choices, as they allow a party to exploit the ambiguity voter experiences when such arguments resonate. Importantly, this would leave the results unchanged.

- $x^{L^*} = -1$ *and* $x^{R^*} = 1$ *if* $a_L \neq v$ *and* $a_R \neq v$
- $x^{L^*} = x^{R^*} = 2\lambda - 1$ *if* $a_L = v$ *and* $a_R = v$*;*
- $x^{R^*} = 1$ *and* $x^{L^*} \in \mathbb{R}$ *if* $a_L = v$ *and* $a_R \neq v$*, and* a_R *resonates;*
- $x^{L^*} = -1$ *and* $x^{R^*} \in \mathbb{R}$ *if* $a_R = v$ *and* $a_L \neq v$*, and* a_L *resonates;*
- $x^{R^*} = 1$ *and* $x^{L^*} = \widetilde{x}^L = -1$ *if* $a_L = v$ *and* $a_R \neq v$*, and* a_R *does not resonate;*
- $x^{L^*} = -1$ *and* $x^{R^*} = \widetilde{x}^R, = 1$ *if* $a_R = v$ *and* $a_L \neq v$*, and* a_L *does not resonate.*

Proof. The first bullet point follows immediately from the assumed constraints, and the second follows from the baseline analysis. Consider the third and fourth bulletpoints, i.e., the case in which one party, j, presents a non-vacuous argument, $-j$ presents a vacuous one, and j's argument resonates. In this case, j is committed to the voter's preferred policy (given her posterior beliefs). Therefore, $-j$ can never win with a platform different from x_j^*. Since our parties are purely policy-motivated, any x_{-j} can trivially be sustained in equilibrium. Finally, consider the fifth bulletpoint (a symmetric reasoning applies to the last one). Here, L makes a vacuous argument, R a non-vacuous one, and R's argument fails to resonate. Recall that $\widetilde{x}^L = -1 > -3$. Thus, L can propose its ideal policy and win with certainty (since $-\widehat{\pi}(-1-1)^2 < -\widehat{\pi}(\widetilde{x}^L - 1)^2$). Thus, in equilibrium, $x^{L^*} = \widetilde{x}^L = -1$. □

Lemma 2F. *Denote with* $\mathcal{V}^i(a^R, a^L)$ *i's continuation value given arguments* a^R *and* a^L*. In equilibrium, we have:*

$$\mathcal{V}^i(a^R = r, a^L = r) = \mathcal{V}^i(a^R = s, a^L = s) \tag{21}$$
$$= -\frac{(1-\pi)}{2}(\tilde{x}^i - \tilde{x}^{-i})^2 - \pi\lambda(-1-\tilde{x}^i)^2 - \pi(1-\lambda)(1-\tilde{x}^i)^2;$$

$$\mathcal{V}^i(a^R = r, a^L = s) = \mathcal{V}^i(a^R = r, a^L = v) = \mathcal{V}^i(a^R = v, a^L = s) \tag{22}$$
$$= -\pi\lambda(-1-\tilde{x}^i)^2 - (1-\pi\lambda)(1-\tilde{x}^i)^2;$$

$$\mathcal{V}^i(a^R = s, a^L = r) = \mathcal{V}^i(a^R = s, a^L = v) = \mathcal{V}^i(a^R = v, a^L = r) \tag{23}$$
$$= -\pi(1-\lambda)(1-\tilde{x}^i)^2 - \Big(1 - \pi(1-\lambda)\Big)(-1-\tilde{x}^i)^2;$$

$$\mathcal{V}^i(a^R = v, a^L = v) = -(1 - 2\lambda - \tilde{x}^i)^2. \tag{24}$$

Proof. Follows directly from Lemma 1F, the voter's optimal electoral strategy, and $\widetilde{x}^R = 1 = -\widetilde{x}^L$. □

Proposition 6F. *The equilibrium of the argumentation stage is as characterized in Proposition 1.*

Proof. The continuation values characterized in Lemma 2F are identical to those characterized in the baseline model in the proof of Lemma 2. Thus, this version of the game is isomorphic to the baseline, and the equilibrium of the argumentation stage must be as characterized in Proposition 1. □

G Proofs for Extension with Multiple Parties

The perspective that we adopt in modeling multiple parties is that there exist two groups on each side of the ideological spectrum, with aligned ideologies but potentially different issue priorities, which may coalesce into different parties. Our model is sufficiently flexible to accommodate any configuration of these groups in the party systems: each group may form its own party, the two groups on each side may coalesce into two main parties, or groups on one side may coalesce while the others remain fragmented.

Policymaking in a multiparty system is inherently more complex than in two-party systems. While fully modeling this process is beyond the scope of this paper, it is reasonable to assume that policymaking outcomes are influenced by the preferences of the electorate on dimensions they care about: as the electorate shifts to the right on a particular dimension, the policy outcome on that dimension tends to move rightward, and vice versa. We capture this idea in a reduced form: instead of explicitly modeling an electoral and policymaking stage as we do in our two-party analysis, we model the parties' strategic behavior at the argumentation stage only, with the parties' utilities determined directly by the voter's posteriors under assumptions that are consistent with the expectation stated in the previous paragraph. Specifically, following the parties' argumentation, the voter updates her beliefs and preferences, as in the baseline model; these preferences, then, determine the realization of the players' payoffs directly. Group g's utility is then defined as follows:

$$u_j^g = \begin{cases} -w_j^g(x_j^v - \tilde{x}_j^g)^2 \quad if \quad \widehat{\pi}_j > 0, \\ -\frac{1}{2}w_j^g(\tilde{x}_j^{R^j} - \tilde{x}_j^g)^2 - \frac{1}{2}w_j^g(\tilde{x}_j^{L^j} - \tilde{x}_j^g)^2 \quad if \quad \widehat{\pi}_j = 0. \end{cases} \tag{25}$$

where w_j^g is the intensity with which group g cares about dimension j. Recall that $\tilde{x}_j^g$ is g's ideal policy on dimension j. In line with our baseline model, we assume that left-wing groups have ideal points at -1 on each dimension, and right-wing groups have ideal points at $+1$. x_j^v is the voter's induced preference on dimension j. $\widehat{\pi}_j$ is the voter's posterior probability that dimension j is relevant for her. If the voter becomes unconcerned with dimension j, we assume that each camp is equally likely to impose its views in the policymaking process. (We could relax this assumption without much effect on the results.)

We consider three possible configurations of the party system, depending on whether groups on either side of the spectrum choose to form independent parties or rather coalesce. If the two groups on the same side of the spectrum coalesce in the same party, the party chooses an argumentation strategy that maximizes the groups' joint welfare. To avoid trivial results, we will assume parties pay an arbitrarily small cost when presenting a non-vacuous argument.

Proposition 7G. *Regardless of the configuration of the party system, there exists a unique $\overline{\pi}_j$ s.t.*

- *if $\pi_i < \overline{\pi}_j$, then in any equilibrium, one party from each side of the spectrum presents a refuting argument on j;*
- *if $\pi_i > \overline{\pi}_j$, then in any Pareto-undominated equilibrium, all parties present a vacuous argument on j.*

$\overline{\pi}_j$ is as characterized in Proposition 1.

Proof. First, notice that no redundant non-vacuous arguments can be sustained in equilibrium, since the parties pay an arbitrary small cost of presenting one. Thus, we must only consider five (classes of) equilibrium candidates for each dimension $j \in \{A, B\}$. (1) One party presents a supporting argument, the others present a vacuous one; (2) One party presents a refuting argument, the others present a vacuous one; (3) One party from each side of the spectrum presents a refuting argument, the others present a vacuous one; (4) One party from each side of the spectrum presents a supporting argument, the others present a vacuous one; (5) All parties present a vacuous argument.

Recall that, given our payoff specification in (25), each group's payoff from any given pair of arguments presented on dimension j in this version of the game coincides exactly with the *equilibrium* continuation value emerging from the baseline model (net of the cost of presenting a non-vacuous argument, which is assumed to be arbitrarily small). This implies that there exists no equilibrium in classes 1, 2 or 4. Equilibria in class 3 always exist, and equilibria in class 5 exist and are Pareto efficient when $\pi_j > \overline{\pi}_j$, where $\overline{\pi}_j$ is as characterized in Proposition 1. □