

Ideology For The Future

Federica Izzo

Current draft: March 28, 2019

Abstract

Do ideologically motivated parties have strategic incentives to lose? I present a model of repeated spatial elections in which the voters face uncertainty about their preferred policy and learn via experience upon observing their payoff realization. The amount of voter learning, I show, depends on the location of the implemented policy: the more extreme the policy is, the more information is generated. This, in turn, creates a trade-off for a party whose ideological stance is unpopular with the electorate, between winning the upcoming election so as to secure policy influence, and changing the voters' preferences so as to win with a better platform in the future. Under some conditions the party gambles on the future: chooses to lose today, in order to change voters' views and win big tomorrow.

Introduction

Barry Goldwater obtained the Republican party presidential nomination in 1964, despite the widespread belief that he was ideologically too extreme to win the general election. Goldwater himself revealed that he never actually thought he could win (Goldwater 1988: 154). Indeed, he went on to lose by a landslide against Lyndon Johnson. Jeremy Corbyn represents a more recent instance of the ‘Goldwater phenomenon’ (Wildavsky 1965). Corbyn won the Labour primaries in 2015 with a 40% margin. Yet, the general opinion was that his leadership would condemn the party to electoral irrelevance (Toynbee 2015). Corbyn’s supporters were aware of his low electoral viability, but were ‘keener on picking a leader who shared their views, rather than someone who was likely to lead Labour to victory’ (YouGov 2015). Indeed, ‘Labour’s new manifesto is the most left-wing since 1983’, when the party ran on a platform labelled as ‘the longest suicide note in history’ (Castle 2017).

These and other examples suggest that political parties sometimes *choose* to settle for electoral defeat: they adopt unpopular positions, even if this means losing the upcoming election for sure. From a rational choice perspective, this is quite puzzling. Extant models of elections predict that instrumentally rational parties will always do whatever it takes to win. Even if a party is motivated solely by ideology, it would never accept a certain electoral defeat. Other authors instead argue that political parties may be willing to lose, but work under the assumption that their members have expressive rather than strategic motivations and care about ideological purity (Aldrich 1983, Wildavsky 1965, Roemer 2001. See also discussion in Strom 1999, Budge et. al 2010).

In this paper, I instead show that ideologically motivated parties may choose to lose for entirely strategic reasons, without any concern for purity. A party whose ideology is unpopular with the electorate is faced with a crucial trade-off, between compromising in order to win the upcoming elections, and changing the voters’ preferences so as to be able to win with a better platform in the future. Under some conditions, the party gambles on the future: chooses to lose today in order to change voters’ views and win big tomorrow.

This paper analyses this trade off within a model of repeated spatial elections with two time periods. The players are two policy motivated parties and a representative voter. In each period, the parties credibly commit to a policy platform along the real line. The voter then decides whom

to elect. The model has two key features. First, the voter faces uncertainty about the exact location of her ideal policy. For example, the voter may not know which policy is most likely to produce her preferred outcome. Thus, we can think about her uncertainty as referring to the true state of the world, representing the policy-outcome mapping. Secondly, the players have different priors on the state of the world but agree to disagree, i.e. they do not update on each other's beliefs. I think about prior beliefs as representing a person's convictions and world views. Thus, while the players are aware of the fact that their priors differ, they do not infer anything from the existence of this disagreement. As a consequence, the voter may only learn via experience: she updates her beliefs upon observing the realization of her first-period payoff. This is a function of the implemented policy, the true state, and of a random shock which complicates the voter's inference problem.

A consequence of this technology is that the amount of voter learning depends on the policy implemented in the first period. Specifically, the voter learns more about the state of the world (and thus her ideal policy) when extreme platforms are enacted. As the policy becomes more extreme, the distance in the expected outcome as a function of the true state increases. As a consequence, each outcome is more informative. In more substantive terms, if the voter likes (dislikes) the outcome of an extreme policy, the policy is likely (unlikely) to be in line with the true state. Conversely, because the voter only observes a noisy signal, the outcome of a moderate policy is much less informative.

Let's now consider the incentives faced by the two parties. The second period is equivalent to a one-shot Downsian model: in equilibrium the parties always converge on the voter's preferred policy. Not so much in the first period. The party whose ideological stance is ex-ante unpopular faces a trade-off, between securing policy influence and forcing the voter to experiment. Suppose that the voter's prior is such that her ex-ante preferred policy is a right-wing one, and consider the problem faced by the left-wing party. The party always has incentives to converge towards the voter's preferences, in order to win the upcoming election and move the implemented platform closer to its own bliss point. This is the usual centripetal tendency that arises in Downsian models. However, the unpopular party also has an incentive to increase the amount of voter learning, in hopes of changing the voter's future policy preferences and being able to implement a better platform tomorrow. The problem the unpopular party faces is that it cannot achieve both goals at once.

This is a direct consequence of the voter's 'bias' against the party. Given the voter's prior, for any pair of policies that leave her indifferent in the first period, the right-wing one is always

further away from zero. Thus, the popular right-wing party can win with relatively more extreme platforms, that would generate a larger amount of information. This creates the trade-off for the unpopular party. It may move slightly closer to the voter and win, thus minimizing the immediate policy losses. However, this would imply that a more moderate policy is implemented and less information is generated. The voter is unlikely to change her mind, and the party will probably have to compromise on a right-wing platform again tomorrow. Conversely, if the unpopular party allows its opponent to win with an extreme right-wing policy, the amount of voter learning increases. If the voter learns that such policy is not aligned with the true state of the world, the unpopular party will be able to win with a left-wing platform in the future.

In other words, the unpopular party must choose between compromising in order to minimize immediate losses – but this means having to compromise again tomorrow – and going all-in hoping to be able to win with a better platform in the future. If the incentives to force the voter to experiment are sufficiently strong, the unpopular party chooses to gamble on the future: lose today to win big tomorrow. This paper characterizes the conditions under which this occurs in equilibrium.

Crucially, I show that extreme policy preferences are not enough for an instrumentally rational party to choose to lose. The ‘gambling’ equilibria can be sustained only if both parties are sufficiently ideological in their prior beliefs, i.e. sufficiently confident that the true state of the world is line with their own policy preferences. Intuitively, the unpopular party is willing to throw out the election only when it believes the gamble is likely to be successful. However, this is not enough. In a Downsian setting, ‘it takes two to gamble’: the popular party must also be willing to increase the amount of voter learning. The popular party has a lot to lose from generating additional information. If it is not sufficiently confident that this will move the voter even closer to its own preferences, the popular party is not willing to take up the gamble and the first period has a unique equilibrium in convergence. Thus, open conflict of (ideological) beliefs is a crucial part of the story.

The nature of electoral competition in this model is very different from the dynamics typically emerging in spatial elections. Probabilistic voting models (e.g. Calvert 1985, Wittman 1987, Groseclose 2001) analyze a trade-off analogous to the one discussed above: policy-motivated parties may adopt a platform that decreases their probability of winning (although they would never accept to lose for sure). However, an instrumental desire to win office still defines the nature of electoral competition. Thus, comparative statics show both parties’ equilibrium platforms always moving

in the same direction of the median voter’s (expected) bliss point. If this ideal policy moves right both platforms move right, with the unpopular party always ‘chasing after’ the voter. Conversely, in the ‘gambling’ equilibria described above the unpopular party’s strategic behavior is driven by the desire to change the voter’s future preferences. As the voter’s right-wing bias increases, the unpopular left-wing party has more to gain and less to lose from forcing her to experiment. Thus, as the voter’s (ex-ante) preferences move to the right, the unpopular party may be willing to go further and further to the left. This allows its opponent to win with a more extreme right-wing platform, thus ensuring that even more information is generated. Therefore, we can – and do, as I discuss below – observe empirical patterns that are consistent with the theory presented here, but are hard to reconcile with probabilistic voting models.

Literature Review

This paper presents a model of repeated spatial elections in which the voter faces uncertainty about her ideal policy. While several works analyse elections under policy-relevant uncertainty, the focus is typically on strategic communication. Politicians have privileged information about the state of the world, and engage in a signalling game with the electorate.¹ The Maskin and Tirole’s (2004) and Canes-Wrone, Herron and Shotts’ (2001) pandering models are obvious examples. Kartik et al. (2015) extends the analysis considering pandering in a Downsian setting. Similarly, in Roemer (1994) voters are uncertain of the functioning of the economy, and fully informed parties compete on policy platforms and on theories of the world.

In this paper, I adopt a different perspective. I consider a setting in which the state of the world (the voter’s ideal policy) is unknown to all players and, as a consequence, the voter may only learn via experience. She updates her beliefs about her ideal platform upon observing the outcome of the policy implemented in the first period. A consequence of this technology is that the amount of voter learning is a function of the location of the implemented policy. This creates incentives for political parties to engage in information control. Thus, when choosing their electoral platforms, parties consider how the policy that is implemented today influences the amount of

¹Kartik, Van Weelden and Wolton (2017) provide an exception. The model features no asymmetry of information at the electoral stage, but the elected politician will discover the true state of the world once in office. This induces the parties to commit to ambiguous platforms in equilibrium.

information the voter will receive tomorrow. Office holders in Dewan and Hortala-Vallve (2015) and Majumdar and Mukand (2004) make similar considerations. However, both papers present variants of the principal agent model in which the incumbent is free to choose his preferred level of policy experimentation. Conversely, I focus on a Downsian setting in which ‘it takes two to gamble’: a gamble takes place in equilibrium only if both parties are willing to generate information. Additionally, in the extant literature policy outcomes reveal information about the office holder’s competence. In the model presented here the voter is instead forced to experiment in order to discover her true policy preferences.

In this perspective, this paper is most closely related to recent work by Callander (2011) and Hirsch (2016). Callander (2011) analyses a spatial election model in which players face uncertainty over the policy-outcome mapping, and update their beliefs upon observing the outcome of the implemented policy. The author assumes that the players know the slope of the policy mapping function (the state of the world), but must learn about the realization of the variance for each policy location, i.e. the exact consequences of each specific policy. As a consequence, small incremental policy changes reveal more information. In contrast, the model presented here adopts a different framework to study policy experimentation (see also Izzo, 2018). Within this framework, uncertainty is over the fundamental underlying state of the world, and the voters’ inference problem is complicated by the presence of a random shock. As such, it is *extreme* policies that reveal more information. Further, focusing on the statically optimal choice for a policy maker, Callander (2011) assumes myopic parties. The main contribution of this paper is instead to investigate how dynamic considerations – i.e. the desire to change voters’ future beliefs – influence the parties’ platform choice. Such dynamic considerations also emerge in Hirsch (2016). The author presents a principal-agent model, in which players have heterogeneous priors about the state of the world. The principal repeatedly chooses a policy, and the agent decides how much effort to exert in its implementation. Effort increases the probability that a policy tailored to the state of the world is successful, but is wasted on a ‘wrong’ policy. Under some conditions, the principal will choose a policy that she considers likely to be wrong, but that the agent believes to be correct, in order to elicit ‘wasted’ effort and eliminate the belief disagreement. In the model that I present below, the unpopular party makes a similar reasoning: it may choose to incur a loss today, in order to generate more information and win big tomorrow.

Finally, a recent working paper by Eguia and Giovannoni (2018) presents an argument analogous to the one advanced here. An office motivated party that experiences a valence disadvantage may choose a radical policy today, in order to acquire ‘ownership’ on that platform. An exogenous shock to the electorate’s preferences may allow the party to reap the benefits of this ‘tactical extremism’, and win with a higher probability in the future. The two works nicely complement each other: Eguia and Giovannoni (2018) consider office-seeking candidates, while I focus on ideologically motivated parties. Further, while in Eguia and Giovannoni (2018) changes in preferences are driven by an exogenous shock, the main contribution of my paper is to present a model in which these changes are instead driven by learning via experience, and arise endogenously as a consequence of the parties’ strategic behaviour. Additionally, I do not assume any ‘stickiness’ in the platforms across periods.

The Model

The model consists of two periods, with an election in each. The players are two policy motivated parties, L and R , and a representative voter V . Before each election, the two parties (simultaneously) commit to a policy platform along the real line, $x_t^i \in \mathbb{R}$, $\forall i \in \{L, R\}$ and $\forall t \in \{1, 2\}$. The voter decides whom to elect. The winner implements the announced platform (credible commitment).

The voter faces uncertainty about the exact location of her ideal policy x_V . This policy can take one of two values that, for simplicity, I assume to be symmetric around 0: $x_V \in \{\underline{\alpha}, \bar{\alpha}\}$ where $\bar{\alpha} = -\underline{\alpha} \geq 0$. We can think about the voter’s uncertainty as referring to the state of the world, representing for example the shape of the policy-mapping function. In other words, the voter does not know which policy is most likely to produce her preferred outcome.

While the true state (i.e. true value of x_V) is unknown to all players, they hold heterogeneous prior beliefs. Players therefore assign different probabilities γ_i , $\forall i \in \{L, V, R\}$ to the voter’s bliss point taking a positive value. Such heterogeneous priors are common knowledge but players agree to disagree, i.e. they do not update on each other’s beliefs. Because this assumption is an important point of departure from the standard tenets of Bayesian rationality, I discuss it in further depth below.

Given common knowledge of heterogeneous priors, the voter only learns via experience. She

observes how much she liked - or disliked - the first-period policy, and updates by using Bayes rule. Formally, the voter's payoff realization is a noisy signal of the state of the world:

$$U_t^V = -(x_V - x_t)^2 + e_t \quad (1)$$

Where

$$e_t \sim U\left[-\frac{1}{2\psi}, \frac{1}{2\psi}\right]$$

The assumption that the noise is drawn from a uniform distribution substantially simplifies the analysis, but is not necessary for the results.

Finally, parties are policy motivated with quadratic loss utility, and I assume that their preferences are not a function of the state of the world. In other words, I assume ideological dogmatism):

$$U_t^i = -(x_i - x_t)^2 \quad (2)$$

$$\forall i \in \{L, R\}$$

Where $x_L \leq 0 \leq x_R$.

Notice that the parties only care about ideology, i.e. assign no value to holding office per se. I discuss this specific assumption, and the results' robustness to relaxing it, in a separate section.

To sum up, the timing of the game is as follows:

1. Nature determines the value of $x_V \in \{\underline{\alpha}, \bar{\alpha}\}$, (that remains unknown to all players) and of the players' priors γ_L , γ_V and γ_R (that become common knowledge)
2. The two parties simultaneously commit to a policy platform $x_1^i \in \mathbb{R}$, $\forall i \in \{L, R\}$
3. The voter decides whom to elect
4. The winner implements the announced platform
5. The voter's first-period payoffs realize
6. Second-period elections are held, as above

7. Second-period payoffs realize, and the game ends

To avoid trivialities, I will assume that the voter's preferred policy is always between the two parties' per-period bliss points, irrespective of her beliefs: $x_L \leq \underline{\alpha} \leq 0 \leq \bar{\alpha} \leq x_R$.

Let me emphasize that the voter has no private information. As a consequence, given any pair of platforms, the parties face no uncertainty over the electoral outcome in the current period. However, uncertainty – and, due to heterogeneous priors, disagreement – exist over what the voter will learn upon observing the first period policy outcome.

Finally, notice that while the model considers parties as unitary actors, it also admits a less literal interpretation. In line with the motivational examples, the game can be interpreted as a reduced-form version of a citizens candidates model with a primary stage. By choosing the candidate, the activists would effectively set the party's electoral platform. Thus, the model speaks to a recurrent argument in the literature, according to which primaries represent a polarizing force and ideologically extreme activists are often unwilling to compromise (Aldrich 1983, Coleman 1971, Brady 2007, Hall 2015). Alternatively, the party's equilibrium platform may be the result of a bargaining process between different factions (as in Levy 2004). This interpretation would be in line with the argument that extreme ideological factions within political parties may put a veto on moderate platforms, even if this means losing for sure (Roemer 2001, Budge et al., 2010).

Heterogeneous Priors and Beliefs as Ideology

Before delving into equilibrium analysis, it is important to discuss in more depth the key assumption that underpins the results: players hold heterogeneous priors on the state of the world, and 'agree to disagree' (Aumann 1976). This represents a departure from canonical models based on the common priors assumption, i.e. the assumption that heterogeneous beliefs can only be due to of information asymmetries. As a consequence, if a conflict of beliefs becomes common knowledge, it is immediately resolved: individuals revise their own priors according to those held by others, and eventually reach full mutual agreement.

In this paper I adopt a different perspective, thinking about prior beliefs as a person's 'mental models, institutions or world views' (Van den Steen 2011: 887). Thus, 'individuals may simply be endowed with different prior beliefs (just as they may be endowed with different preferences)' (Che

and Kartik 2009). In a similar vein, Callander argues that ‘much political disagreements is over beliefs (...), that we may think of as ideology’ (2011: 657). Hafer and Landa (2005, 2007) also see ideology and beliefs as closely connected, thinking of a player’s ideology as the likelihood of being persuaded by a left-wing argument versus a right-wing one. Analogous intuitions are presented by Piketty (1995), Benabou and Tirole (2006) and McMurray (2016).

In line with these arguments, I model parties’ beliefs as a second dimension of their ideology: each party is convinced that the true state of the world is aligned with its own policy preferences. The left (right) wing party always wants to implement a left (right) wing policy, irrespective of the state of the world. However, the party also believes that such policy is in line with the true state. Formally, I assume that $\gamma_L = 1 - \gamma_R = \epsilon$, where ϵ takes an arbitrarily small value. I will then show that the results can be sustained under less restrictive conditions, as long as both parties are sufficiently ideological in their beliefs.

Conceptualizing priors as ideology, I allow open conflicts of beliefs to be sustained in equilibrium. Players have different ‘world views’ that translate into different beliefs about the true state. Simply becoming aware of the existence of this conflict is not enough to solve it. Indeed, quite the opposite. ‘Individuals with belief conflicts think that they can persuade each other by taking actions that will produce more information, each expecting it to prove that they were right’ (Hirsch, 2016: 70).

In addition to the scholars mentioned above, several others have allowed players to ‘agree to disagree’ (see Yildiz 2004, Smith and Stam 2004, Minozzi 2013, Ashworth and Sasso 2017). Thus, while somewhat unorthodox, this approach is not unprecedented in the literature.

Analysis: Learning

The voter’s learning plays a crucial role in the mechanism the model identifies. Thus, before analyzing the player’s equilibrium behavior it is important to understand how learning occurs.

The voter’s first-period payoff realization is a noisy signal of the state of the world. In other words, the voter considers how much she liked or disliked the first-period policy, and updates her beliefs by using Bayes’ rule. The analysis reveals a crucial feature of the learning process: the amount of information received by the voter depends on the location of the policy implemented in the first period. Specifically, the voter learns more about the state of the world (i.e. the location

of her ideal policy) when more extreme platforms are enacted. As the implemented policy moves away from zero, the distance in the expected outcome as a function of the true state increases. As a consequence, each signal is more informative. In more substantive terms, if the voter likes (dislikes) the outcome of an extreme policy, it is likely that such policy is (is not) in line with her true preferences. However, the outcome of a moderate policy is much less informative. It is harder for the voter to understand whether the policy produced a good outcome because it is in line with the true state, or despite this not being true but due to the presence of a small shock.

This feature emerges in a very stark form in a world in which the noise e_t is uniformly distributed. Denote as μ_V the voter's posterior that $x_V = \bar{\alpha}$, given her own payoff realization U_1^V , the first-period policy x_1 and her prior γ_V . The following Lemma holds:

Lemma 1. *The voter learning satisfies the following properties:*

- (i) *Her posterior μ_V takes one of three values: $\mu_V \in \{0, \gamma, 1\}$;*
- (ii) *The more extreme the policy implemented in the first period x_1 , the higher the probability that $\mu_V \neq \gamma$;*
- (iii) *There exists a policy x' such that if $|x_1| \geq |x'|$, then $\mu_V \neq \gamma$ with probability 1.*

Lemma 1 tells us that upon observing her first-period payoff realization, the voter learns either everything or nothing about the state of the world. The more extreme the implemented policy, the more likely it is to generate an informative signal. While a formal proof of this Lemma is presented in the Appendix, the underlying reasoning is easy to illustrate graphically.

In Figure 1, the solid lines represent the voter's expected payoff as a function of the implemented policy x_1 , for the two possible values of x_V . Thus, the thick increasing solid curve is $-(x_1 - \bar{\alpha})^2$ and the thin decreasing solid curve is $-(x_1 - \underline{\alpha})^2$. For any policy different from zero, the voter's expected payoff is always different in the two states of the world. However, recall that the actual payoff realization is also a function of the realization of the shock e_1 . The dashed curves represent therefore the maximum and minimum possible values of the payoff realization, once we take the shock into account. Suppose that the true state is positive ($x_V = \bar{\alpha}$). Then, for any policy x_1 the actual payoff realization can fall *anywhere* on the line between the two thick increasing dashed curves (representing, respectively, $-(x_1 - \bar{\alpha})^2 + \frac{1}{2\psi}$ and $-(x_1 - \bar{\alpha})^2 - \frac{1}{2\psi}$). Analogously, if the true state is negative the payoff realization can be anywhere on the line between the thin decreasing




Figure 1: Voter's payoff realization as a function of first-period policy. The thick (thin) curves represent the case in which $x_V = \bar{\alpha}$ ($x_V = \alpha$). Solid curves are the voter's expected payoff $E[U_1^v]$, dashed ones represent $E[U_1^v] - \frac{1}{2\psi}$ and $E[U_1^v] + \frac{1}{2\psi}$

dashed curves.

The presence of the shock creates a partial overlap in the support of the payoff realization for a positive and negative state of the world: for any given policy $x_1 \in (-x', x')$, there exist values of the voter's payoff that may be observed whatever the true state. Consider for example policy x , as represented in the graph. Any payoff realization falling between the gray and black bullets may be observed with positive probability under both states of the world. Straightforwardly, if the payoff realization falls outside this range of overlap, it constitutes a fully informative signal. There is only one state of the world that could have generated that specific realization: the voter simply likes the policy too much, or too little, for this to be justified as a consequence of the shock. Thus, upon observing her payoff, the voter learns the true state (i.e. discovers the true value of x_V). Conversely, any payoff realization that falls inside the range of overlap is completely uninformative. Due to the assumption that the shock is uniformly distributed, any such realization has exactly the same probability of being observed under the two states of the world. Thus, the voter learns nothing and must go back to her prior beliefs. The more extreme the implemented policy, the closer the gray and black bullets get, the smaller the range of overlap and the higher the probability that the voter will learn the true state.

Let me emphasize that my results only require that extreme policies are more informative than moderate ones. The assumption that the noise is drawn from a uniform distribution is not necessary

to generate this result. Consider for example a world in which the noise is normally distributed with full support. The learning process would be much smoother: any signal would be somewhat informative, but never fully so. However, it would still be the case that extreme policies generate more information. As the implemented policy becomes more extreme, the distance in the expected outcomes as a function of the state increases. This, in turn, increases the signal’s informativeness.²

The Voter

In what follows, I will assume without loss of generality that the voter’s prior is ‘biased’ in favor of the right-wing party, so that her ex-ante preferred policy is a positive one: $\gamma_V > \frac{1}{2}$.³ Thus, I refer to the left-wing (right-wing) party as the unpopular one (popular one). In order to simplify the presentation of the results, but without much loss of substance, I assume that $\bar{\alpha} \leq x'$, i.e. even under complete information the voter’s preferred policy is never sufficiently extreme to guarantee learning with probability one. For ease of presentation I initially consider a myopic voter. I then show that the (qualitative) results are robust to assuming a forward looking, and fully patient, voter.

Let us focus first on the voter’s strategy. Her equilibrium behavior is straightforward:

Lemma 2. *In each period, the voter elects the party whose platform is closer to her preferred policy (given her own beliefs).*

The voter’s preferred policy in the first period is a function of her prior: $\bar{\alpha}(2\gamma_V - 1)$. In the second period it will instead reflect her updated beliefs: $\bar{\alpha}(2\mu_V - 1)$ (where, given Lemma 1, $\mu_V \in \{0, \gamma_V, 1\}$). The proof of this Lemma follows the usual argument and is therefore omitted.

The Parties

Consider now the parties’ platform choice. Absent any future concerns, the second-period subgame is exactly equivalent to a one-shot Downsian game. Thus, the following Lemma holds:

²Indeed, all I need to sustain this result is that the noise distribution satisfies the monotonic likelihood ratio property.

³The results hold symmetrically for $\gamma_V < \frac{1}{2}$. The strict inequality is necessary.

Lemma 3. *The second-period subgame has a unique equilibrium, in which both parties commit to the voter’s preferred policy: $x_2^{L*} = x_2^{R*} = \bar{\alpha}(2\mu_V - 1)$*

The proof follows the usual argument. Divergent platforms can never be sustained in equilibrium in the second period. If neither of the two parties is at the voter’s bliss point, at least one of them can always increase its payoff by moving closer to the voter and winning for sure. If only one of the two parties is at the voter’s bliss point, it can always deviate to a winning platform that strictly increases its own payoff. Suppose instead that the parties converge on the voter’s preferred platform. Neither of them can change the policy implemented in equilibrium by unilateral deviation. Therefore, convergence on the voter’s preferences can always be sustained in equilibrium.

It is easy to see that the second result can be extended to the first period: the game always has an equilibrium in which the parties converge on the voter’s preferred policy in both periods. However, the key argument of this paper is that this classic equilibrium is not always unique and does not always capture the nature of electoral competition. In what follows, I will show that the unpopular party’s strategic behavior is sometimes driven by the incentives to gamble on the future and change the voter’s preferences, even at the cost of losing for sure.

The Parties’ Utility

Lemma 1 shows that the location of the policy implemented in the first period has a crucial impact on the voter learning. The more extreme the policy is, the larger the variance in the distribution of her posterior beliefs (i.e. the larger the likelihood that $\mu_V \neq \gamma_V$). The voter’s posterior in turns determines the platform that will be enacted in the second period (Lemma 3). Thus, the policy implemented in the first period has a twofold effect on the parties’ expected utility. A direct effect on their first-period payoff, and an indirect one on their expected future utility (via the voter learning). The direct effect is clear: each party’s utility decreases as the platform moves away from its per-period bliss point. Figure 2 represents the left-wing party’s first-period payoff. Straightforwardly, as the policy moves to the right away from x_L , the party’s utility strictly decreases. The indirect effect is more subtle. Each party is convinced that the true state of the world is in line with its own policy preferences (i.e. $\gamma_L = 1 - \gamma_R = \epsilon$, where ϵ takes an arbitrarily small value). Thus, each believes that information would always move the voter’s future preferences closer to its own. As




Figure 2: Party L 's first-period utility as a function of the implemented policy

consequence, each party's expected future utility increases as the policy implemented in the first period becomes more extreme, *both to the left and to the right of 0*. Recall that this expectation is the 'subjective' one, as a function of the party's own prior.

The overall impact of the first-period policy on the parties' expected utility will depend on the combination of the direct and indirect effects. Focus again on the unpopular left-wing party (with symmetric results holding for the right-wing one). If we consider a left wing policy ($x_1 < 0$) moving to the right away from x_L , direct and indirect effect go in the same direction. The party's immediate payoff decreases, and as the policy moves closer to zero it also (weakly) reduces the amount of voter learning. This also implies that the policy that maximizes the party's expected utility – which I denote as x_L^g – is (weakly) to the left of x_L . Conversely, when a positive policy moves further to the right, direct and indirect effect have different signs. As the policy moves to the right the party's first-period payoff decreases. At the same time, however, a more extreme policy being implemented implies that the voter is more likely to learn the true state of the world, which increases the party's expected future utility. If the indirect effect is sufficiently strong, the party's expected utility has a second (local) maximum in the positive numbers, which I denote as x_L^{pos} . The following Lemma holds:

Lemma 4. *There exist unique $\tilde{\alpha}$ and \tilde{x}_L such that if $\bar{\alpha} > \tilde{\alpha}$ and $x_L < \tilde{x}_L$, then L 's expected*




Figure 3: Party L 's expected utility as a function of first-period policy

utility on $[0, \infty]$ is non monotonic with a maximum at $x_L^{pos} > 0$. Otherwise, L 's expected utility is monotonically decreasing on $[0, \infty]$.

The indirect effect is stronger if information has a large impact on the voter's policy preferences: the larger $\bar{\alpha}$, the higher the expected gain from increasing the amount of voter learning. Additionally, the more extreme the party is, the more it benefits from moving the voter's future preferences to the left (given concave utility). Thus, if the conditions in Lemma 4 are satisfied the indirect effect dominates, and the left-wing party's overall utility increases as the implemented policy moves further to the right in the range $[0, x_L^{pos}]$ (as depicted in Figure 3).⁴ In what follows, I show that the presence of this non-monotonicity is what allows gambling behavior to emerge in equilibrium.

Gambling on the Future

Let's now focus on the incentives the parties face in the first period platform game. Consider the popular party R . Recall that (by assumption) $x_R > \bar{\alpha}$, where x_R is the party's 'static' bliss point (i.e. the policy that maximises its utility in the current period). Additionally, since the party's expected future utility is increasing in the amount of voter learning, its welfare maximizing policy

⁴Notice that, since the probability of learning is not smooth in x_1 , neither is the utility function: it kinks at $-x'$, 0 and x' (see Lemma 1).




Figure 4: Platforms symmetric around the voter’s preferred policy

x_R^g is (weakly) more extreme than x_R . This implies that, in equilibrium, the winning platform must always be to the right of the voter’s preferred policy ($\bar{\alpha}(2\gamma_V - 1)$). Given any policy to the left of this point, the right-wing party can always find a different platform that increases both its own and the voter’s payoff. In particular, for any policy $x < 0$, the party can move to the symmetric $-x > 0$. This guarantees the same amount of learning, but increases both the voter’s and the party’s immediate payoff. The popular right-wing party would therefore never allow its opponent to win with a policy to the left of the voter.

Should the same reasoning apply to the left-wing party, the usual Downsian dynamics would emerge, thereby leading to a unique equilibrium in full convergence. Instead, the unpopular party faces a trade off between securing policy influence and forcing the voter to experiment. This is a direct consequence of the voter’s ‘bias’ against the party. Given $\gamma_V > \frac{1}{2}$, for any pair of platforms that leave the voter indifferent, the right-wing one is always further away from zero (Figure 4). Thus, the popular party can win with relatively more extreme platforms, that would therefore generate a larger amount of information.

The unpopular party must choose between compromising today so as to move the implemented platform closer to its preferred policy, and allowing its opponent to win in order to increase the amount of voter learning. The party always has an incentive to converge towards the voter’s preferred platform, so as to win the upcoming election and move the implemented policy to the left. However, this would imply that little information is generated, the voter is unlikely to change her beliefs, and the party will have to compromise on a right-wing platform again tomorrow. Conversely, if the party allows its opponent to win with an extreme right-wing policy, the probability that the voter learns the true state increases and the party is more likely to be able to win with a left-wing platform in the future.

If the incentives to force the voter to experiment are sufficiently strong, the unpopular party gambles on the future: allows the right-wing opponent to win, in the hope that the voter will learn that its policies are not aligned with the true state. The unpopular party chooses to lose today in

order to change voters' views and win big tomorrow. In what follows, I establish the conditions under which this behavior can be sustained in equilibrium.

I denote a *gambling equilibrium* an equilibrium of the game in which, in the first period:

(i) the parties adopt platforms on opposite sides of the voter's preferred policy:

$$x_1^{L*} < \bar{\alpha}(2\gamma_V - 1) < x_1^{R*};$$

(ii) the unpopular party L loses with probability 1.

Notice that any equilibrium satisfying (i) must also meet condition (ii). As mentioned above, the popular party would never allow its opponent to win with a policy to the left of the voter. Thus, any divergence equilibrium must be a gambling equilibrium.

Proposition 1 identifies necessary and sufficient conditions for gambling equilibria to exist. Proposition 2 then characterizes the range of platforms that can be sustained in a gambling equilibrium.

Proposition 1. *The exist unique $\widehat{x}_L \in (0, \widetilde{x}_L)$ and $\widetilde{\alpha}$ such that gambling equilibria exist if and only if:*

- *The unpopular party is sufficiently extreme: $x_L < \widehat{x}_L$*
- *Learning the true state has a sufficiently large impact on the voter's preferences: $\bar{\alpha} > \widetilde{\alpha}$*

The thresholds are a function of the other parameters in the model. The conditions ensure that $x_L^{pos} > \bar{\alpha}(2\gamma_V - 1)$, i.e. L 's expected utility is increasing in x_1 at $x_1 = \bar{\alpha}(2\gamma_V - 1)$ (see Figure 5).⁵ Substantively, the expected gain from increasing the amount of voter learning is sufficiently large that the unpopular party is willing to throw out the first-period election.

The qualitative conditions are in line with those identified in Lemma 4 (indeed, the condition on $\bar{\alpha}$ is identical). If the voter receives no additional information, the parties will converge on $\bar{\alpha}(2\gamma_V - 1)$ in the second period. Suppose instead that the voter learns that the true state of the world is in line with the left-wing party's ideology; then, the second-period equilibrium policy will move to $\underline{\alpha}$. Straightforwardly, the gain from a successful gamble is therefore increasing in $\bar{\alpha} = -\underline{\alpha}$. Additionally, the value of moving tomorrow's equilibrium policy increases as the party's bliss point

⁵Recall that x_L^{pos} is the maximum of the left-wing party's expected utility in the positive numbers (Lemma 4).




Figure 5: Players' utility as a function of first-period policy. The solid line represents the left-wing party's expected utility in the whole game, while the dashed one represents the voter's expected utility in the first period.

x_L moves to the left. The unpopular party is willing to gamble only when its ideological preferences are sufficiently extreme.

Further, Corollary 1 shows that gambling equilibria are 'more likely' to exist the larger γ_V : the stronger the voter's right wing 'bias', the easier it is to satisfy the conditions in Proposition 1. The incentives to force the voter to experiment are stronger the further away her preferences would be from the party's, should she receive no additional information. As γ_V increases, the voter's initial preferences move further to the right, and the gain from a successful gamble increases. In other words, the less popular the party is to begin with, the less it has to lose and the more to gain from changing the voter's future preferences.

Corollary 1. *The likelihood that gambling equilibria exist (in the sense of set inclusion) increases as the voter's right-wing bias gets stronger (i.e. $\frac{\partial \hat{x}_L}{\partial \gamma_V} > 0$ and $\frac{\partial \hat{\alpha}}{\partial \gamma_V} < 0$)*

Finally, Proposition 2 identifies the range of platforms that can be sustained in a gambling equilibrium. For ease of presentation, the proposition is derived under the assumption that $x_R > x'$, where x' is the smallest (positive) policy that guarantees learning with probability 1 (see Lemma 1). The assumption simply ensures that $x_R^g > x_L^{pos}$, where x_R^g is the right-wing party's first-period preferred policy. The assumption will be relaxed in Proposition 4.

Proposition 2. *There exists a unique $x_L^{Min}(\bar{\alpha}, \gamma_V, x_L) \geq 2\bar{\alpha}(2\gamma_V - 1) - x_L^{pos}$ such that in any gambling equilibrium, platforms satisfy:*

1. $x_1^{R*} - \bar{\alpha}(2\gamma_V - 1) = \bar{\alpha}(2\gamma_V - 1) - x_1^{L*}$;
2. $x_1^{L*} \geq x_L^{Min}$

Point 1 indicates that in any gambling equilibrium the two parties must be adopting platforms equidistant from the voter's preferred policy. The proof is straightforward: for any pair of asymmetric policies at least one of the parties can deviate to a winning platform that strictly increases its own expected utility. If $x_1^{R*} \neq x_R^g$, R can always find a winning platform closer to x_R^g . If $x_1^{R*} = x_R^g$, the left-wing party can move to x_L^{pos} and win, while strictly increasing its expected utility. Point 2 then identifies the range of platforms that can be sustained in a gambling equilibrium. Straightforwardly, the unpopular party would never allow its opponent to win with a policy to the right of x_L^{pos} . The lower bound of the range is therefore always (weakly) larger than the symmetric $2\bar{\alpha}(2\gamma_V - 1) - x_L^{pos}$.

Notice that, in equilibrium, the voter must be breaking indifference in favour of the popular party R . With any other indifference breaking rule R has a profitable deviation to move slightly closer to the voter and increase its probability of winning. The conjectured equilibria collapse and the parties are driven all the way to full convergence. Thus, in a gambling equilibrium the unpopular party is choosing to lose the election with probability one, even if an arbitrarily small deviation would be enough to win for sure. When instead the conditions in Proposition 1 are not satisfied, electoral competition is driven by the parties' desire to minimize immediate losses and the classic Downsian results hold. The game has a unique equilibrium, in which the parties converge on the voter's bliss point in both periods.⁶

The above results show that, under some conditions, the nature of electoral competition may instead be very different from the classic dynamics emerging in spatial models. While probabilistic voting models analyze a trade-off analogous to the one presented in this paper, electoral competition is still driven by the parties' (instrumental) desire to win office. As a consequence, comparative statics show both equilibrium platforms always moving in the same direction as the (expected) median voter. If the voter moves right, both parties move right in equilibrium. The unpopular party is therefore always chasing after the voter

⁶If the conditions are satisfied there exist other equilibria, in which both parties adopt the same platform in the range $[\bar{\alpha}(2\gamma_V - 1), 2\bar{\alpha}(2\gamma_V - 1) - x_L^{Min}]$, where x_L^{Min} is as defined in Proposition 2.

Conversely, in a gambling equilibrium electoral competition is driven by the unpopular party’s desire to move the electorate’s future preferences closer to its own, even at the cost of losing for sure. As the voter’s right-wing bias increases, the unpopular party has more to gain and less to lose from forcing her to experiment. The party may therefore be willing to go further and further to the left, thus allowing its opponent to win with a more and more extreme right-wing platform that further increases the amount of voter learning. The following Corollary holds:

Corollary 2. *There exists a $\underline{x}_L < \widehat{x}_L$ such that if $x_L > \underline{x}_L$, then $\frac{\partial x_L^{Min}}{\partial \gamma_V} < 0$: as the voter’s right wing bias increases, the unpopular party is willing to move further to the left in equilibrium.*

This result indicates that we may observe empirical patterns that would allow us to adjudicate between competing explanations. Indeed, recent work by Margalit et. al (2017) presents evidence that is hard to reconcile with probabilistic voting models, and is instead consistent with Corollary 2. The authors analyze data from OECD countries since the post-war period, and find that parties tend to move away from the center following an electoral loss. ‘Under standard Downsian logic, parties should move towards the median voter in the electorate (...). If a loss implies that a party was too far away from the median, then the predicted reaction should be a shift to the center’ (p. 4). The model presented here provides a potential explanation as for why a different pattern instead emerges in the data. Learning that the electorate is further to the right increases the unpopular left-wing party’s incentives to gamble, potentially inducing it to move its electoral platform further to the left away from the (median) voter .

In concluding this section, it is important to discuss the impact of a specific assumption used here: parties care only about policy and derive no benefit from holding office per se. To simplify the presentation of the results, the model analysed in this paper maintains several of the key features of the standard spatial model. In particular, the two parties must move simultaneously, and the left-wing (right-wing) party can credibly commit even to extreme right-wing (left-wing) platforms. These assumptions are quite restrictive but they usually bear no impact on the equilibrium results. Not so much in this model. Indeed, in the current set-up gambling equilibria exist only if parties are purely policy motivated. However, relaxing either one, or both, of these assumptions would allow gambling behaviour to emerge in equilibrium even if parties care about office as well as policy. Suppose for example that the two parties have full commitment ability, but can choose the timing

of their platform announcement. Then, gambling equilibria survive as long as office rents are not too large. This is due to the fact that in a gambling equilibrium both parties must have incentives to generate information. Further, allowing for sequential moves would also refine our equilibrium predictions. In particular, in any gambling equilibrium of the sequential moves game platforms are as follows: $x_1^{R*} = \bar{\alpha}(2\gamma_V - 1) - x_L^{Min}$ and $x_1^{L*} - \bar{\alpha}(2\gamma_V - 1) = \bar{\alpha}(2\gamma_V - 1) - x_1^{R*}$, where x_L^{Min} is as defined in Proposition 2.

Alternatively, we could assume that the parties must move at the same time but are somewhat limited in their commitment ability. For example, Levy (2004) speculates that an internal bargaining process between competing factions is what sustains the credibility of electoral promises. Thus, parties can only credibly commit to policies in the Pareto set of the party's members. Alternatively, it may be argued that individual politicians have no credible commitment ability, therefore a party can only propose a platform if it is the true bliss point of one of its members (Krasa and Polborn, 2018). There may exist some overlap in the credible sets of the two parties. Crucially, both may be able to commit to the voter's ideal policy. Nonetheless, as long as the right-most (left-most) platform that the left-wing (right-wing) party can promise is not too extreme, gambling equilibria survive for sufficiently low office rents.

Parties' Beliefs and Ideology

I have so far assumed that each party assigns probability (arbitrarily close to) 1 to the true state of the world being in line with its own ideology, i.e. each believes information would *always* move the voter's future preferences closer to its own. However, gambling equilibria survive under less restrictive conditions. Propositions 3 and 4 generalize the results presented in the previous section, without imposing any prior assumption on the parties' beliefs.

Proposition 3. *There exist unique $\tilde{\alpha}$, x_L^\dagger , x_R^\dagger and $\gamma < \gamma_V$ such that gambling equilibria exist if and only if:*

- *Learning the true state has a sufficiently large impact on the voter's preferences: $\bar{\alpha} > \tilde{\alpha}$*
- *The parties are sufficiently extreme: $x_L < x_L^\dagger$ and $x_R > x_R^\dagger$*
- *The parties are sufficiently ideological in their beliefs: $\gamma_L < \gamma < \gamma_R$*

The thresholds are a function of the other parameters in the model. The first condition is exactly as in Proposition 1. Even when it recognizes that information may move the voter to the right (i.e. $\gamma_L > 0$), the unpopular left-wing party is willing to gamble only if the stakes are sufficiently high. If the voter learns that the true state is right-wing, her second-period policy preferences move to $\bar{\alpha}$. Therefore, as $\bar{\alpha} = -\underline{\alpha}$ increases a failed gamble becomes more and more costly. However, at the same time the gain from a successful gamble also increases – moving the voter all the way to $\underline{\alpha}$ – and to a larger extent (given $\gamma_V > \frac{1}{2}$). Thus, learning the true state must have a sufficiently large impact on the voter’s preferences. Additionally, both parties must be sufficiently extreme in their preferences and ideological in their beliefs. Intuitively, the unpopular party is willing to lose the first period election only if it believes the gamble is likely to be successful. Thus, L must be sufficiently confident that the true state is in line with its own preferences: γ_L must be sufficiently low. However, this is not enough. In a Downsian setting ‘it takes two to gamble’: the popular party must also be willing to increase the amount of voter learning. The right-wing party is ready to take the bet only if it believes information is likely to move the voter even closer to its own bliss point: γ_R must be sufficiently high. In the conjectured equilibria the popular party is winning with probability 1, and implementing a right-wing platform. It is not straightforward to see why it may have a profitable deviation. However, R has a lot to lose from forcing the voter to experiment, especially when its ideological stances are very popular to begin with (i.e. the voter’s prior is high). If γ_R is too low, the party has an incentive to prevent information generation, and the conjectured equilibria collapse.

Further, notice that $\gamma_V > \gamma$. This implies that gambling equilibria can be sustained when the voter and the right-wing party have exactly the same beliefs ($\gamma_R = \gamma_V$), or when the two parties’ priors are arbitrarily close ($\gamma_L = \gamma - \varepsilon$ and $\gamma_R = \gamma_V + \varepsilon$, where ε takes an arbitrarily small value). However, a disagreement between the voter and the unpopular party is always necessary. In other words, the unpopular party must always hold ideological beliefs. Interestingly, the higher the stakes, the smaller the minimum disagreement required to sustain gambling in equilibrium (i.e. $\gamma_V - \gamma$ is decreasing in $\bar{\alpha}$).

These results show that ideological beliefs are a crucial part of the story. Extreme preferences are not enough for an instrumentally rational party to be willing to throw out an election. The party must also be convinced that its ideology is in line with the state of the world. Thus, ideological

‘extremism’ in both beliefs and policy preferences is necessary for gambling behavior to emerge in equilibrium. However, the analysis also reveals that extreme beliefs may to a certain extent substitute for extreme preferences. Specifically, the following comparative statics hold:

Corollary 3. *As the parties become more ideological in their beliefs, gambling equilibria can be sustained under more and more moderate policy preferences: $\frac{\partial x_L^\dagger}{\partial \gamma_L} > 0$ and $\frac{\partial x_R^\dagger}{\partial \gamma_R} < 0$*

The intuition is clear: the more ideological a party is in its beliefs, the more it expects to gain from forcing the voter to experiment. As a consequence, the party will be willing to gamble under relatively less extreme policy preferences.

Finally, Proposition 4 identifies the range of platforms that can be sustained in a gambling equilibrium. Before stating the Proposition, let me introduce some useful notation. Denote as \tilde{x} the policy that maximises R ’s expected utility in the range $[0, x']$. If $x_R < x'$, then \tilde{x} is the right-wing party’s welfare maximising policy (i.e. $\tilde{x} = x_R^g$).⁷ Conversely, if $x_R > x'$ the right-wing party’s expected utility has a first maximum at \tilde{x} and a second one at x_R . Depending on the parameter values, either \tilde{x} or x_R is the function’s global maximum. Notice that, if the conditions in Proposition 3 hold, \tilde{x} is always to the right of the voter’s preferred point (the conditions guarantee that R ’s expected utility is increasing in at $\bar{\alpha}(2\gamma_V - 1)$).

The following holds.

Proposition 4. *Suppose that $\tilde{x} \geq x_L^{pos}$. Then, there exists a unique $\widehat{x_L^{Min}} \geq 2\bar{\alpha}(2\gamma_V - 1) - x_L^{pos}$ such that in any gambling equilibrium platforms satisfy:*

1. $x_1^{R*} - \bar{\alpha}(2\gamma_V - 1) = \bar{\alpha}(2\gamma_V - 1) - x_1^{L*}$
2. $x_1^{L*} \geq \widehat{x_L^{Min}}$

Suppose instead that $\tilde{x} < x_L^{pos}$. Then, there exists a unique $\widetilde{x_L^{Min}} \geq 2\bar{\alpha}(2\gamma_V - 1) - \tilde{x}$ such that any pair of platforms satisfying:

1. $x_1^{R*} - \bar{\alpha}(2\gamma_V - 1) = \bar{\alpha}(2\gamma_V - 1) - x_1^{L*}$
2. $x_1^{L*} \geq \widetilde{x_L^{Min}}$

⁷Recall that x_R is the right-wing party’s preferred policy in a one shot game, i.e. absent learning. x' is the smallest positive policy that guarantees learning with probability one.

can be sustained in a gambling equilibrium. Further, if $E[U_L(2\bar{\alpha}(2\gamma_V - 1) - \tilde{x})] \leq E[U_L(\tilde{x})]$ then there exist also asymmetric gambling equilibria in which $x_1^{R*} = \tilde{x}$ and $x_1^{L*} < 2\bar{\alpha}(2\gamma_V - 1) - x_1^{R*}$. No other gambling equilibrium exists.

First, consider the case in which $\tilde{x} \geq x_L^{pos}$, i.e. the right-wing party's expected utility is increasing at x_L^{pos} . In this case, the equilibrium correspondence has the same properties as identified in Proposition 2. The left-wing party would never allow its opponent to win with a policy to the right of x_L^{pos} . As such, the left-most platform that can be sustained in equilibrium is weakly larger than $2\bar{\alpha}(2\gamma_V - 1) - x_L^{pos}$. Further, the two parties must always adopt symmetric policies. For any pair of asymmetric platforms, R could always deviate to a winning policy that strictly increases its expected utility (i.e. closer to \tilde{x}).

Suppose instead that $\tilde{x} < x_L^{pos}$. In this case, the right-wing party is never willing to commit to x_L^{pos} . It could always deviate to \tilde{x} and strictly increase both its own and the voter's payoff. Indeed (given the definition of \tilde{x}) the same reasoning applies to any platform in $[\tilde{x}, x']$. Further, recall that $x_L^{pos} \leq x'$ therefore no platform to the right of x' can ever be sustained in equilibrium. As such, in any gambling equilibrium $x_1^{R*} \leq \tilde{x}$. Straightforwardly, in any equilibrium in which $x_1^{R*} < \tilde{x}$, the two parties must be adopting symmetric platforms. The right-wing party can otherwise always find a winning policy that strictly increases its expected utility. Conjecture now an asymmetric gambling equilibrium in which the right-wing party proposes \tilde{x} , and the left-wing party commits to a policy x_1^L further from the voter's bliss point. Such an equilibrium can never be sustained if the left-wing party can move slightly to the right of $2\bar{\alpha}(2\gamma_V - 1) - \tilde{x}$ and strictly increase its expected utility. If instead $E[U_L(2\bar{\alpha}(2\gamma_V - 1) - \tilde{x})] \leq E[U_L(\tilde{x})]$, the unpopular party can do nothing better than allow its opponent to win (recall that L 's expected utility is monotonically decreasing on $[2\bar{\alpha}(2\gamma_V - 1) - \tilde{x}, 0]$ and monotonically increasing on $[0, \tilde{x}]$). The conjectured equilibrium can be sustained for any x_1^L if \tilde{x} is the right-wing party's utility global maximum (i.e. $x_R < x'$), and for a sufficiently moderate x_1^L otherwise.

A Look at a Forward Looking Voter

I have so far worked under the assumption that the voter is myopic, and fully discounts the future. While there are substantive reasons to defend such an assumption, it is important to highlight that

the results survive with a forward looking, and fully patient, voter. In this section I analyze the model presented above, but allow the voter to have a positive discount factor $\delta > 0$.

Proposition 5. *There exist unique $\tilde{\alpha}$, γ , $\ddagger x_L \leq \dagger x_L$ and $\ddagger x_R \geq \dagger x_R$ such that gambling equilibria exist if and only if the following conditions are satisfied:*

- *Learning the true state has a sufficiently large impact on the voter's preferences: $\bar{\alpha} > \hat{\alpha}$*
- *The parties are sufficiently ideological in their beliefs: $\gamma_L < \gamma < \gamma_R$*
- *The parties are sufficiently extreme: $x_L < x_L^\ddagger$ and $x_R > x_R^\ddagger$*

The conditions guarantee that the parties' expected utility is increasing at $x_1 = x_V^g$, where x_V^g is the forward looking voter's preferred policy in period one (Figure 6). This is (analogously to what established in the previous sections) necessary and sufficient for gambling equilibria to exist. The qualitative results are as in Proposition 3: gambling equilibria exist if and only if information has a sufficiently large impact on the voter's future preferences, and the parties are sufficiently extreme in both their ideological preferences and ideological beliefs. However – while the conditions on $\bar{\alpha}$, γ_R , and γ_L are exactly the same as in Proposition 3 – those on the parties' preferences are a function of the voter's discount factor δ . The more patient the voter is, the more extreme the parties need to be for gambling equilibria to exist (i.e. x_L^\ddagger is decreasing in δ and x_R^\ddagger is increasing in δ). The forward looking voter's expected utility is increasing in the probability of learning. As a consequence, x_V^g is always more extreme than $\bar{\alpha}(2\gamma_V - 1)$. As δ increases, the voter's desire to learn the true state gets stronger, and her preferred policy moves further to the right. For gambling behavior to be sustained in equilibrium the parties must be more and more extreme, ensuring that they have an incentive to further increase the amount of voter learning.

Characterizing the full range of platforms that can be sustained in a gambling equilibrium is more challenging than when considering a myopic voter. This is due to the fact that a forward looking voter's expected utility may not be single peaked. Indeed, if the value of information is sufficiently large, the voter's expected utility will have a second (local) maximum in the negative numbers (denoted as x_V^{neg} in Figure 6). Thus, for any platform $x > x_V^g$ there may exist multiple negative policies that leave the voter weakly better off. This makes it hard to identify pairs of platforms such that the left-wing party has no profitable deviation.




Figure 6: Forward looking voter's expected utility as a function of first-period policy

However, there must always exist a range of positive policies that provide the voter with strictly higher utility than x_V^{neg} . In particular, there always exist a pair of policies $\underline{x} \in [0, x_V^g)$ and $\bar{x} > x_V^g$ such that $E[U_V(\underline{x})] = E[U_V(\bar{x})] = E[U_V(x_V^{neg})]$, and $E[U_V(x)] > E[U_V(x_V^{neg})]$ for any $x \in (\underline{x}, \bar{x})$ (see Figure 6). The existence of this range allows us to partially characterize the equilibrium correspondence.

Proposition 6. *Any pair of platforms satisfying:*

1. $E[U_V(x_1^{L*})] = E[U_V(x_1^{R*})]$
2. $\underline{x} \leq x_1^{L*} \leq x_V^g \leq x_1^{R*} \leq \min \in \{\bar{x}, x_L^{pos}, \tilde{x}\}$, where \tilde{x} is the maximum of R 's expected utility in the range $[0, x']$

can be sustained in a a gambling equilibrium.

Conclusion

Political parties sometimes adopt extreme positions, even if this comes at the expenses of their electoral success. This behavior is puzzling from a rational choice perspective, and is usually ascribed to ideological dogmatism and expressive concerns for ideological purity. In this paper, I

have shown that ideologically motivated parties may instead choose to lose for entirely strategic reasons. A party whose ideology is unpopular with the electorate faces a trade off, between securing immediate policy influence and changing the voter’s future preferences. If the party is sufficiently extreme and ideological in its beliefs, it may adopt the ‘strategy of changing preferences of voters, so that when it wins at some future date, it can be with a better policy’ (Roemer, 2001: 154). The unpopular party chooses to lose today and gamble on the future: allows its opponent to win with an extreme policy that increases the amount of voter learning. If the gamble is successful, and the voter learns that she dislikes the opponent’s policies, the ex-ante unpopular party will be able to win with a better platform in the future.

As it is often the case with research papers, in completing this project I have encountered a trade-off between presenting the most stylized model that would allow me to explore and easily present the key dynamic trade-off of interest, and introducing a richer setting to delve deeper into the nature of uncertainty and the voters’ learning process. Here, I have decided in favour of the first alternative. In particular, the model presented in this paper embeds the assumption that the world is simple, with few unknowns (here, an unknown state of the world representing the slope of the policy mapping function). In principle, then, policy making is easy: in each period, the decision maker only has to think about how likely it is that the state of the world is a right-wing versus a left-wing one. However, learning is made difficult by the fact that policy outcomes are noisy (i.e., realized outcomes are a function of random shocks). At the other extreme, we find models such as Callander (2011) based on the assumption that there is no noise, but the world itself is complex with many (or even infinite) unknowns. Policy making is hard, and today’s outcome influences tomorrow’s decision indirectly via learning and directly by impacting voter’s willingness to settle or keep experimenting. Thus, adopting one or the other conceptualization of uncertainty bears important consequences on the voters’ learning process, the nature of policy making and the dynamic link between periods. Callander assumes that players know the slope of the policy mapping function (which in my setting would represent the state of the world). However, if we amend his framework to allow for some degree of uncertainty, then the key trade-off highlighted here should continue to hold (with extreme policies revealing more information about the correct state). Nonetheless, it would be interesting to analyse whether further insights emerge when we consider this richer setting, as well as whether the conditions for the emergence of gambling equilibria are qualitatively similar. This remains an

objective for future research.

References

- Aldrich, J. H. (1983), 'A Downsian spatial model with party activism', *American Political Science Review*, 77: 4, 974-990.
- Ashworth, S., & Sasso, G. (2018), 'Delegation to an Overconfident Expert'.
- Aumann, R. J. (1976), 'Agreeing to disagree', *The annals of statistics*, 1236-1239.
- Benabou, R., & Tirole, J. (2005), 'Belief in a just world and redistributive politics', (No. w11208). *National Bureau of Economic Research*.
- Budge, I., Ezrow, L., & McDonald, M. D. (2010), 'Ideology, party factionalism and policy change: An integrated dynamic theory' *British Journal of Political Science*, 40: 4, 781-804.
- Callander, S. (2011), 'Searching for Good Policies', *American Political Science Review*, 105:4, 643-662.
- Calvert, R. (1985), 'Robustness of the Multidimensional Voting Model: Candidates Motivations, Uncertainty and Convergence', *American Journal of Political Science*, 29: 69-95.
- Canes-Wrone, B., Herron, M. C., & Shotts, K. W. (2001), 'Leadership and pandering: A theory of executive policymaking', *American Journal of Political Science*, 532-550.
- Castle, S. (2017), 'Labour Party in Britain Approves Jeremy Corbyn's Sharp Left Turn', *The New York Times* <https://www.nytimes.com/2017/05/11/world/europe/uk-labour-leaked-manifesto-jeremy-corbyn.html>
- Che, Y. K., & Kartik, N. (2009), 'Opinions as incentives', *Journal of Political Economy*, 117: 5, 815-860.
- Dewan, T., & Hortala-Vallve, R. (2015), 'Teams of Rivals: Learning in the Cabinet and its Shadow.', Working Paper.
- Downs, A. (1957), *An Economic Theory of Democracy*, Harper and Row, New York.
- Eguia, J. X., & Giovannoni, F. (2018), 'Tactical Extremism'.

- Goldwater, BM. (1988), *Goldwater*, Doubleday, New York.
- Groseclose, Tim (2001) 'A model of candidate location when one candidate has a valence advantage' *American Journal of Political Science*, 45: 4, 826-866.
- Hafer, C., & Landa, D. (2007), 'Deliberation as self-discovery and institutions for political speech', *Journal of Theoretical Politics*, 19: 3, 329-360.
- Hafer, C., & Landa, D. (2005), 'Deliberation as Self-Discovery and Group Polarization', New York University.
- Hirsch, A. V. (2016), 'Experimentation and persuasion in political organizations', *American Political Science Review*, 110: 01, 68-84.
- Hotelling, H. (1929), 'Stability in Competition', *The Economic Journal*, 39: 153, 41-57.
- Izzo, F. (2018), 'With Friends Like These, Who Needs Enemies?', Working Paper.
- Kartik, N., Squintani, F., & Tinn, K. (2015), 'Information revelation and pandering in elections', Columbia University, New York, 36.
- Kartik, N., Van Weelden, R., & Wolton, S. (2016), 'Electoral Ambiguity and Political Representation.' *Available at SSRN 2682691*.
- Majumdar, S., & Mukand, S. W. (2004), 'Policy gambles.', *The American Economic Review*, 94: 4, 1207-1222.
- Margalit, Y., Slough T. & Ting M. (2018), 'After Defeat: Governing Party Response to Electoral Loss'.
- Maskin, E., & Tirole, J. (2004), 'The politician and the judge: Accountability in government', *American Economic Review*, 94: 4, 1034-1054.
- McMurray, J. (2017), 'Ideology as opinion: a spatial model of common-value elections', *American Economic Journal: Microeconomics*, 9: 4, 108-40.
- Minozzi, W. (2013), 'Endogenous Beliefs in Models of Politics', *American Journal of Political Science*, 57: 3, 566-581.

- Muller, W. C., & Strom, K. (Eds.). (1999), *Policy, office, or votes?: how political parties in Western Europe make hard decisions*, Cambridge University Press.
- Piketty, T. (1995), 'Social mobility and redistributive politics.' *The Quarterly journal of economics*, 551-584.
- Roemer, J. (2001), *Political Competition. Theory and Applications*, Harvard University Press, Harvard.
- Roemer, J. (1994), 'The Strategic Role of Party Ideology When Voters Are Uncertain about How the Economy Works', *American Political Science Review*, 88: 02, 327-335.
- Smith, Alastair, & Allan C. Stam (2004), 'Bargaining and the Nature of War', *Journal of Conflict Resolution*, 48: 12, 783-813.
- Toynbee, P. (2015), 'Free to dream, I'd be left of Jeremy Corbyn. But we can't gamble the future on him', *The Guardian* <http://www.theguardian.com/commentisfree/2015/aug/04/jeremy-corbyn-gamble-labour-future-yvette-cooper-best-chance>.
- Van den Steen, E. (2011), 'Overconfidence by Bayesian-rational agents', *Management Science*, 57: 5, 884-896.
- Wildavsky, A. (1965), 'The Goldwater phenomenon: Purists, politicians, and the two-party system', *The Review of Politics*, 27: 3, 386-413.
- Wittman, D. (1983), 'Candidate Motivations: A Synthesis of Alternative Theories?', *American Political Science Review*, 77: 1, 152-157.
- Yildiz, M. (2004), 'Waiting to Persuade', *The Quarterly Journal of Economics*, 119: 2, 223-248.

Appendix

Lemma 1: *voter learning satisfies the following properties:*

- (i) Her posterior μ_V takes one of three values: $\mu_V \in \{0, \gamma, 1\}$;
- (ii) The more extreme the policy implemented in the first period x_1 , the higher the probability that $\mu \neq \gamma$;
- (iii) There exists a policy x' such that if $|x_1| \geq |x'|$, then $\mu_V \neq \gamma$ with probability 1.

Proof. The proof of Claims 1 and 2 below is necessary and sufficient to prove Lemma 1.

Claim 1: *Let $x_t \geq 0$.*

(i) *A payoff realization $U_t^v \notin [-(x_t - \bar{\alpha})^2 - \frac{1}{2\psi}, -(x_t - \underline{\alpha})^2 + \frac{1}{2\psi}]$ is fully informative. Upon observing $U_t^v > -(x_t - \underline{\alpha})^2 + \frac{1}{2\psi}$, the players form posterior beliefs that $x_V = \bar{\alpha}$ with probability 1. Similarly, upon observing $U_t^v < -(x_t - \bar{\alpha})^2 - \frac{1}{2\psi}$ the players form beliefs that $x_V = \underline{\alpha}$ with probability 1.*

(ii) *A payoff realization $U_t^v \in [-(x_t - \bar{\alpha})^2 - \frac{1}{2\psi}, -(x_t - \underline{\alpha})^2 + \frac{1}{2\psi}]$, is uninformative. Upon observing U_t^v , players confirm their prior belief that $x_V = \bar{\alpha}$ with probability γ_i , $\forall i \in \{R, V, L\}$.*

Symmetric results apply when $x_t < 0$.

Proof. The proof of part (i) is trivial given the boundedness of the distribution of e , and is therefore omitted. Part (ii) follows straightforwardly from applying Bayes rule. Recall that the voter's payoff realization U_t^v is a function of the implemented policy (x_t) the voter's true bliss point (x_V) and the noise term (e): $U_t^v = -(x_V - x_t)^2 + e$. Denote as $f(\cdot)$ the PDF of e . Then,

$$\text{prob}(x_V = \bar{\alpha} | U_t^v) = \frac{f(U_t^v + (x_t - \bar{\alpha})^2)\gamma}{f(U_t^v + (x_t - \bar{\alpha})^2)\gamma + f(U_t^v + (x_t - \underline{\alpha})^2)(1 - \gamma)} \quad (3)$$

Given the assumption that ϵ is uniformly distributed

$$f(U_t^v + (x_t - \bar{\alpha})^2) = f(U_t^v + (x_t - \underline{\alpha})^2) \quad (4)$$

Therefore the above simplifies to

$$\text{prob}(x_V = \bar{\alpha}|U_t^v) = \gamma \quad (5)$$

This concludes the proof of Claim 1. \square

Claim 1 proves that players either observe an uninformative or a fully informative signal. Claim 2 shows that the policy choice determines the expected probability that the signal will be informative. The more extreme the implemented policy, the higher such probability.

Claim 2: Let L be a binary indicator, taking value 1 if the players learn the true value of x_V at the end of period 1, and 0 otherwise. There exists $x' = \frac{1}{4\bar{\alpha}\psi}$ such that

- For all $|x_1| \geq |x'|$

$$\text{Prob}(L = 1|x_1) = 1 \quad (6)$$

- For all $x_1 \in [0, x')$

$$\text{Prob}(L = 1|x' \geq x_1 \geq 0) = 4\bar{\alpha}\psi x_1 \quad (7)$$

- For all $x_1 \in (-x', 0]$

$$\text{Prob}(L = 1| -x' \leq x_1 \leq 0) = -4\bar{\alpha}\psi x_1 \quad (8)$$

Proof. Let me first prove the existence of point x' . From Claim 1, x' is the point such that for any policy $|x| \geq |x'|$, the interval $[-(x_t - \bar{\alpha})^2 - \frac{1}{2\psi}, -(x_t - \underline{\alpha})^2 + \frac{1}{2\psi}]$ is empty. This requires

$$-(x_t - \underline{\alpha})^2 + \frac{1}{2\psi} + (x_t - \bar{\alpha})^2 + \frac{1}{2\psi} \leq 0 \quad (9)$$

Recall that $\bar{\alpha} = -\underline{\alpha}$, thus the above reduces to

$$x \geq \frac{1}{4\bar{\alpha}\psi} = x' \quad (10)$$

To complete the proof, assume $x_1 \in (0, x')$. The expected probability of the realized outcome being informative is:

$$\begin{aligned}
 & Prob(L = 1 | \gamma, 0 < x_1 < x') = \\
 & \gamma [Prob(-(x_t - \bar{\alpha})^2 + e_1 > -(x_t - \underline{\alpha})^2 + \frac{1}{2\psi})] + (1 - \gamma) [Prob(-(x_t - \underline{\alpha})^2 + e_1 < -(x_t - \bar{\alpha})^2 - \frac{1}{2\psi})] \quad (11)
 \end{aligned}$$

Given the symmetry

$$Prob(-(x_t - \bar{\alpha})^2 + e_1 > -(x_t - \underline{\alpha})^2 + \frac{1}{2\psi}) = Prob(-(x_t - \underline{\alpha})^2 + e_1 < -(x_t - \bar{\alpha})^2 - \frac{1}{2\psi}) \quad (12)$$

(15) simplifies to

$$Prob(L = 1 | x_1 > 0) = Prob(-(x_t - \bar{\alpha})^2 + e_1 > -(x_t - \underline{\alpha})^2 + \frac{1}{2\psi}) = 4\bar{\alpha}\psi x_1 \quad (13)$$

Similar calculations produce the result for $x_1 \in (-x', 0]$.

This concludes the proof of Claim 2 □

and thus of Lemma 1 □

The Parties' Utility

In this section I will characterize the policies x_L^g and x_L^{pos} (symmetric results apply for the right-wing party), and present the proof of Lemma 4.

Denote as $\beta(x_1)$ the probability of the voter learning the true state of the world (as a function of the policy implemented in the first period). Given $\gamma_L = \epsilon \approx 0$, the left-wing party's (subjective) expected utility can be written as:

$$-(x_1 - x_L)^2 - (1 - \beta(x_1))(\bar{\alpha}(2\gamma_V - 1) - x_L)^2 - \beta(x_1)(\underline{\alpha} - x_L)^2 \quad (14)$$

Notice that the party's utility is increasing in $\beta(x_1)$, given the assumption on γ_L . From Lemma 1 we know that $\beta(x_1)$ is not a smooth function of x_1 : it kinks at $-x'$, 0 and x' . Thus, we must

analyze the utility function piecewise.

Consider first the case in which $x_L \leq -x'$. Then, L 's expected utility as a function of x_1 has the following properties:

- In the range $[-\infty, -x']$ it is concave and non monotonic with global maximum at $x_L^g = x_L$. Every policy in this range guarantees learning with probability 1. Thus, as x_1 moves away from x_L it only has a negative direct effect on the party's payoff.
- In the range $[-x', 0]$ it is strictly decreasing. As the policy moves to the right the party's immediate utility decreases. The probability of the voter learning the true state is also reduced, which implies lower expected future utility
- In the range $[0, x']$ the party faces a trade-off, that is analyzed in more details below.
- In the range $[x', \infty]$ it is strictly decreasing. Every policy in this range guarantees learning with probability 1. Thus, as x_1 moves to the right it only has a negative direct effect on the party's payoff.

Consider now the case in which $x_L > -x'$. Then, L 's expected utility as a function of x_1 has the following properties:

- In the range $[-\infty, -x']$ it is strictly increasing. Every policy in this range guarantees learning with probability 1. Thus, as x_1 moves closer to x_L it only has a positive direct effect on the party's payoff.
- In the range $[-x', 0]$ it is concave and non-monotonic with global maximum at $x_L^g \in [-x', x_L]$.

This is the policy that solves the following maximization problem:

$$\begin{aligned}
 & \underset{x_1}{\text{maximise}} && -(x_1 - x_L)^2 - (1 + 4\bar{\alpha}\psi x_1)(\bar{\alpha}(2\gamma_V - 1) - x_L)^2 + 4\bar{\alpha}\psi x_1(\underline{\alpha} - x_L)^2 \\
 & \text{subject to} && x_1 \in \left[-\frac{1}{4\bar{\alpha}\psi}, 0\right]
 \end{aligned} \tag{15}$$

- In the range $[0, x']$ the party faces a trade-off, that is analyzed in more details below.

- In the range $[x', \infty]$ it is strictly decreasing. Every policy in this range guarantees learning with probability 1. Thus, as x_1 moves away from x_L it only has a negative direct effect on the party's payoff.

Lemma 4: *There exist unique $\tilde{\alpha}$ and \tilde{x}_L such that if $\bar{\alpha} > \tilde{\alpha}$ and $x_L < \tilde{x}_L$ then L 's expected utility on $[0, \infty]$ is non monotonic with a maximum at $x_L^{pos} > 0$. Otherwise, L 's expected utility is monotonically decreasing on $[0, \infty]$.*

Proof. From the discussion above we know that L 's utility is always monotonically decreasing in the range $[x', \infty]$. Conversely, in the range $[0, x']$ the party faces a trade off. As the policy moves to the right the party's immediate payoff decreases, while its future expected payoff increases. The maximization problem is:

$$\begin{aligned} & \underset{x_1}{\text{maximise}} && -(x_1 - x_L)^2 - (1 - 4\bar{\alpha}\psi x_1)(\bar{\alpha}(2\gamma_V - 1) - x_L)^2 - 4\bar{\alpha}\psi x_1(\underline{\alpha} - x_L)^2 \\ & \text{subject to} && x_1 \in [0, \frac{1}{4\bar{\alpha}\psi}] \end{aligned} \tag{16}$$

The solution to this maximisation problem is $x^* = \min \in \{\max \in \{0, x_L - 8\bar{\alpha}^2\psi(x_L\gamma_V + \bar{\alpha}\gamma_V(1 - \gamma_V))\}, \frac{1}{4\bar{\alpha}\psi}\}$. Thus, if $x_L - 8\bar{\alpha}^2\psi(x_L\gamma_V + \bar{\alpha}\gamma_V(1 - \gamma_V)) \leq 0$, the function is monotonically decreasing on $[0, \infty]$. Otherwise, it is non monotonic with maximum at $x_L^{pos} = \min \in \{x_L - 8\bar{\alpha}^2\psi(x_L\gamma_V + \bar{\alpha}\gamma_V(1 - \gamma_V)), \frac{1}{4\bar{\alpha}\psi}\}$. Therefore, the condition for non-monotonicity is $x_L - 8\bar{\alpha}^2\psi(x_L\gamma_V + \bar{\alpha}\gamma_V(1 - \gamma_V)) > 0$. This yields:

$$x_L < \frac{-8\bar{\alpha}^3\psi\gamma_V(1 - \gamma_V)}{8\bar{\alpha}^2\psi\gamma_V - 1} \tag{17}$$

and

$$\bar{\alpha}^2 > \frac{1}{8\psi\gamma_V} \tag{18}$$

□

Proposition 1: *The exist unique $\hat{x}_L \in (0, \tilde{x}_L)$ and $\tilde{\alpha}$ such that Gambling equilibria exist if and only if:*

- The unpopular party is sufficiently extreme: $x_L < \widehat{x}_L$
- Learning the true state has a sufficiently large impact on the voter's preferences: $\bar{\alpha} > \widetilde{\alpha}$

Proof. Necessary and sufficient condition for gambling equilibria to exist is that L 's expected utility is increasing at $x_1 = \bar{\alpha}(2\gamma_V - 1)$, i.e. $x_L^{pos} > \bar{\alpha}(2\gamma_V - 1)$. Notice that $\bar{\alpha}(2\gamma_V - 1) < \frac{1}{4\bar{\alpha}\psi}$ (given the assumption that $\bar{\alpha} < \frac{1}{4\bar{\alpha}\psi}$). Thus, we do not have to worry about the case in which (16) has a corner solution at $\frac{1}{4\bar{\alpha}\psi}$, and the condition is:

$$x_L - 8\bar{\alpha}^2\psi(x_L\gamma_V + \bar{\alpha}\gamma_V(1 - \gamma_V)) > \bar{\alpha}(2\gamma_V - 1) \quad (19)$$

The above can be satisfied if and only if the LHS is decreasing in x_L . Thus, we obtain:

$$x_L < \frac{-\bar{\alpha}(2\gamma_V - 1) - 8\bar{\alpha}^3\psi\gamma_V(1 - \gamma_V)}{8\bar{\alpha}^2\psi\gamma_V - 1} \quad (20)$$

And

$$\bar{\alpha}^2 > \frac{1}{8\psi\gamma_V} \quad (21)$$

The proof of Corollary 1 follows straightforwardly from above. \square

Proposition 2: *There exists a unique $x_L^{Min}(\bar{\alpha}, \gamma_V, x_L) \geq 2\bar{\alpha}(2\gamma_V - 1) - x_L^{pos}$ such that in any gambling equilibrium, platforms satisfy:*

1. $x_1^{R*} - \bar{\alpha}(2\gamma_V - 1) = \bar{\alpha}(2\gamma_V - 1) - x_1^{L*}$;
2. $x_1^{L*} \geq x_L^{Min}$

Proof. The proof for Point 1 is provided in the main body of the paper. Here I provide the proof for point 2. x_L^{Min} is the left-most platform that L is willing to adopt in equilibrium. First of all, notice that $x_L^{Min} \geq 2\bar{\alpha}(2\gamma_V - 1) - x_L^{pos}$: L would never allow its opponent to win with a platform to the right of x_L^{pos} . In other words, for any pair of platforms satisfying Point 1 and such that $x_L^{Min} < 2\bar{\alpha}(2\gamma_V - 1) - x_L^{pos}$, L would have a profitable deviation to move to x_L^{pos} . It follows that $x_L^{Min} = 2\bar{\alpha}(2\gamma_V - 1) - x_L^{pos}$ when $2\bar{\alpha}(2\gamma_V - 1) - x_L^{pos} \geq 0$. Recall, in fact, that the party's utility is monotonically increasing on $[0, x_L^{pos}]$. Suppose instead that $2\bar{\alpha}(2\gamma_V - 1) - x_L^{pos} < 0$. Then, the

following Corollary holds:

Corollary 1A: *Suppose that $2\bar{\alpha}(2\gamma_V - 1) - x_L^{pos} < 0$. Then, $x_L^{Min} = \max \in \{2\bar{\alpha}(2\gamma_V - 1) - x_L^{pos}, \hat{x}\}$, where $\hat{x} \leq 0$ is such that $E[U_L(\hat{x})] = E[U_L(2\bar{\alpha}(2\gamma_V - 1) - \hat{x})]$*

Proof. First of all let me prove the existence of a (unique) policy \hat{x} .

Claim 1. *There exists a unique policy $\hat{x} \leq 0$ such that: (i) $E[U_L(\hat{x})] = E[U_L(2\bar{\alpha}(2\gamma_V - 1) - \hat{x})]$, (ii) for any $x < \hat{x}$, $E[U_L(x)] > E[U_L(2\bar{\alpha}(2\gamma_V - 1) - x)]$ and (iii) for any $\hat{x} < x < 0$, $E[U_L(x)] < E[U_L(2\bar{\alpha}(2\gamma_V - 1) - x)]$.*

Proof. Given $x_L^{pos} > 2\bar{\alpha}(2\gamma_V - 1)$, L 's expected utility is monotonically increasing on $[0, \bar{\alpha}(2\gamma_V - 1)]$. It follows straightforwardly that:

$$E[U_L(x)] < E[U_L(2\bar{\alpha}(2\gamma_V - 1) - x)] \quad (22)$$

When $x = 0$. Additionally, it is easy to see that the following holds:

$$E[U_L(x)] > E[U_L(2\bar{\alpha}(2\gamma_V - 1) - x)] \quad (23)$$

When $x \leq -x'$ (since both x and $2\bar{\alpha}(2\gamma_V - 1) - x$ guarantee learning with probability 1, but x is always closer to x_L).

Thus, there must exist (at least) one policy $\hat{x} \in (-x', 0)$ such that

$$E[U_L(\hat{x})] = E[U_L(2\bar{\alpha}(2\gamma_V - 1) - \hat{x})] \quad (24)$$

The uniqueness of \hat{x} follows straightforwardly from the fact that $E[U_L(2\bar{\alpha}(2\gamma_V - 1) - x)]$ is monotonically decreasing on $[-x', 0]$, while $E[U_L(x)]$ is either monotonically decreasing or concave with maximum at x_L^g (see analysis at p. 35).

□

Claim 1 (along with Point 1 in Proposition 2) implies that $x_L^{Min} \geq \hat{x}$: for any pair of platforms symmetric around the voter and such that $x_1^{L*} < \hat{x}$, L has a profitable deviation to make an arbitrarily small move to the right and win for sure. Further, recall that L 's expected utility

is monotonically decreasing on $[2\bar{\alpha}(2\gamma_V - 1) - x_L^{pos}, 0]^8$ and monotonically increasing on $[0, x_L^{pos}]$. Additionally (as discussed in the main body), notice that x_L^{Min} must always be to the right of $2\bar{\alpha}(2\gamma_V - 1) - x_L^{pos}$. Thus, it follows straightforwardly from Claim 1 that $x_L^{Min} = \max \in \{2\bar{\alpha}(2\gamma_V - 1) - x_L^{pos}, \hat{x}\}$.

This concludes the proof of Corollary 1A □

and Proposition 2. □

Corollary 2: *There exists $\underline{x}_L < \widehat{x}_L$ such that if $x_L > \underline{x}_L$, then $\frac{\partial x_L^{Min}}{\partial \gamma_V} < 0$: as the voter's right wing bias increases, the unpopular party is willing to move further to the left in equilibrium.*

Proof. Suppose that $x_L > \underline{x}_L$, where the condition guarantees that $x_L - 8\bar{\alpha}^2\psi(x_L\gamma_v + \bar{\alpha}\gamma_v(1 - \gamma_v)) < \min\{\frac{1}{4\bar{\alpha}\psi}, 2\bar{\alpha}(\gamma_V - 1)\}$.⁹ Then, $x_L^{Min} = 2\bar{\alpha}(\gamma_V - 1) - x_L - 8\bar{\alpha}^2\psi(x_L\gamma_v + \bar{\alpha}\gamma_v(1 - \gamma_v))$. Thus, x_L^{Min} is decreasing in γ_V iff:

$$4\bar{\alpha} + 8\bar{\alpha}^2\psi(x_L + \bar{\alpha}(1 - 2\gamma_V)) < 0 \quad (25)$$

Which reduces to:

$$x_L < \frac{2\bar{\alpha}^2\psi(2\gamma_V - 1) - 1}{2\bar{\alpha}\psi} \quad (26)$$

From the proof of Proposition 1 we know that gambling equilibria exists if and only if the following condition is satisfied:

$$x_L < \frac{-\bar{\alpha}(2\gamma_V - 1) - 8\bar{\alpha}^3\psi\gamma_V(1 - \gamma_V)}{8\bar{\alpha}^2\psi\gamma_V - 1} \quad (27)$$

It is easy to verify that the RHS in condition (27) is strictly smaller than the RHS in (26). As such, (26) is never binding and $x_L > \underline{x}_L$ is sufficient to guarantee that $\frac{\partial x_L^{Min}}{\partial \gamma_V} < 0$. □

Proposition 3: *There exist unique $\widetilde{\bar{\alpha}}$, x_L^\dagger , x_R^\dagger and $\gamma < \gamma_V$ such that gambling equilibria exist if and only if:*

- *Learning the true state has a sufficiently large impact on the voter's preferences : $\bar{\alpha} > \widetilde{\bar{\alpha}}$*
- *The parties are sufficiently extreme: $x_L < x_L^\dagger$ and $x_R > x_R^\dagger$*

⁸It is straightforward to verify that $2\bar{\alpha}(2\gamma_V - 1) - x_L^{pos}$ is always to the right of the function's maximum on $[-x', 0]$.

⁹When the conditions in Proposition 1 are satisfied, $x_L - 8\bar{\alpha}^2\psi(x_L\gamma_v + \bar{\alpha}\gamma_v(1 - \gamma_v))$ is always decreasing in x_L .

- *The parties are sufficiently ideological in their beliefs: $\gamma_L < \gamma < \gamma_R$*

Proof. As in Proposition 1, necessary condition for the conjectured equilibria to be sustained is that $x_L^{pos} > \bar{\alpha}(2\gamma_V - 1)$:

$$x_L - 8\bar{\alpha}^2\psi(x^L(\gamma_V - \gamma_L) + \bar{\alpha}\gamma_V(1 - \gamma_V)) > \bar{\alpha}(2\gamma_V - 1) \quad (28)$$

The above can be satisfied only if the LHS is decreasing in x_L . Thus we obtain

$$x^L < \frac{-\bar{\alpha}(2\gamma_V - 1) - 8\bar{\alpha}^3\psi\gamma_V(1 - \gamma_V)}{8\bar{\alpha}^2\psi(\gamma_V - \gamma_L) - 1} \quad (29)$$

$$\gamma_L < \gamma_V - \frac{1}{8\bar{\alpha}^2\psi} \quad (30)$$

and

$$\bar{\alpha}^2 > \frac{1}{8\psi\gamma_V} \quad (31)$$

However this is not sufficient. It is also necessary for the right-wing party's utility to be strictly increasing at $x_1 = \bar{\alpha}(2\gamma_V - 1)$.¹⁰ R 's expected utility on $[0, x']$ is:

$$\begin{aligned} E[U_R(x_1)] &= -(x_1 - x_R)^2 - (1 - 4\bar{\alpha}\psi x_1)(\bar{\alpha}(2\gamma_V - 1) - x_R)^2 \\ &\quad - 4\bar{\alpha}\psi x_1[\gamma_R(\bar{\alpha} - x_R)^2 + (1 - \gamma_R)(\underline{\alpha} - x_R)^2] \end{aligned}$$

Thus $\frac{\partial E[U_R(x_1)]}{\partial x_1} = -2(x_1 - x_R) + 4\bar{\alpha}\psi(2\bar{\alpha}(2\gamma_V - 1) - x_R)^2 - 4\bar{\alpha}\psi(\gamma_R(\bar{\alpha} - x_R)^2 + (1 - \gamma_R)(\underline{\alpha} - x_R)^2)$.

The equilibrium condition is therefore:

$$-\bar{\alpha}(2\gamma_V - 1) + x_R + 8\bar{\alpha}^2\psi(x_R(\gamma_R - \gamma_V) - \bar{\alpha}\gamma_V(1 - \gamma_V)) > 0 \quad (32)$$

Which can be rewritten as:

¹⁰Notice that, given $x_R > \bar{\alpha}$, this is always true under the assumption that $\gamma_R \approx 1$, which was used to derive Propositions 1 and 2.

$$x_R > \frac{\bar{\alpha}(2\gamma_V - 1) + 8\bar{\alpha}^3\psi\gamma_V(1 - \gamma_V)}{8\bar{\alpha}^2\psi(\gamma_R - \gamma_V) + 1} \quad (33)$$

Which requires

$$\gamma_R > \gamma_V - \frac{1}{8\bar{\alpha}^2\psi} \quad (34)$$

□

Proposition 4: Suppose that $\tilde{x} \geq x_L^{pos}$. Then, there exists a unique $\widehat{x}_L^{Min} \geq 2\bar{\alpha}(2\gamma_V - 1) - x_L^{pos}$ such that in any gambling equilibrium platforms satisfy:

1. $x_1^{R*} - \bar{\alpha}(2\gamma_V - 1) = \bar{\alpha}(2\gamma_V - 1) - x_1^{L*}$
2. $x_1^{L*} \geq \widehat{x}_L^{Min}$

Suppose instead that $\tilde{x} < x_L^{pos}$. Then, there exists a unique $\widetilde{x}_L^{Min} \geq 2\bar{\alpha}(2\gamma_V - 1) - \tilde{x}$ such that any pair of platforms satisfying:

1. $x_1^{R*} - \bar{\alpha}(2\gamma_V - 1) = \bar{\alpha}(2\gamma_V - 1) - x_1^{L*}$
2. $x_1^{L*} \geq \widetilde{x}_L^{Min}$

can be sustained in a gambling equilibrium. Further, if $E[U_L(2\bar{\alpha}(2\gamma_V - 1) - \tilde{x})] \leq E[U_L(\tilde{x})]$ then there exist also asymmetric gambling equilibria in which $x_1^{R*} = \tilde{x}$ and $x_1^{L*} < 2\bar{\alpha}(2\gamma_V - 1) - x_1^{R*}$. No other gambling equilibrium exists.

Proof. The proof proceeds as for Proposition 2 and is therefore omitted. □

Proposition 5: There exist unique $\widetilde{\alpha}$, γ , $\ddagger x_L \leq \dagger x_L$ and $\ddagger x_R \geq \dagger x_R$ such that gambling equilibria exist if and only if the following conditions are satisfied:

- Learning the true state has a sufficiently large impact on the voter's preferences: $\bar{\alpha} > \widetilde{\alpha}$
- The parties are sufficiently ideological in their beliefs: $\gamma_L < \text{gamma} < \gamma_R$
- The parties are sufficiently extreme: $x_L < x_L^\ddagger(\delta)$ and $x_R > x_R^\ddagger(\delta)$

Proof. First of all we must calculate the voter's optimum x_V^g . This is the policy that solves the following maximization problem:

$$\begin{aligned} & \underset{x_1}{\text{maximise}} && -\gamma_V(x_1 - \bar{\alpha})^2 - (1 - \gamma_V)(x_1 - \underline{\alpha})^2 - \delta(1 - 4\bar{\alpha}\psi x_1)[\gamma_V(\bar{\alpha}(2\gamma_V - 1) - \bar{\alpha})^2 + (1 - \gamma_V)(\bar{\alpha}2\gamma_V)^2] \\ & \text{subject to} && x_1 \leq \frac{1}{4\bar{\alpha}\psi} \end{aligned} \quad (35)$$

$x_V^g = \min\{\frac{1}{4\bar{\alpha}\psi}, \bar{\alpha}(2\gamma_V - 1) + 8\delta\bar{\alpha}^3\psi\gamma_V(1 - \gamma_V)\}$. Given $\bar{\alpha} < x' = \frac{1}{4\bar{\alpha}\psi}$, $x_V^g = \bar{\alpha}(2\gamma_V - 1) + 8\delta\bar{\alpha}^3\psi\gamma_V(1 - \gamma_V)$. Thus, necessary and sufficient conditions for the existence of gambling equilibria are:

$$x_L - 8\bar{\alpha}^2\psi(x^L(\gamma_V - \gamma_L) + \bar{\alpha}\gamma_V(1 - \gamma_V)) - \bar{\alpha}(2\gamma_V - 1) - 8\delta\bar{\alpha}^3\psi\gamma_V(1 - \gamma_V) > 0 \quad (36)$$

And

$$x_R + 8\bar{\alpha}^2\psi(x_R(\gamma_R - \gamma_V) - \bar{\alpha}\gamma_V(1 - \gamma_V)) - \bar{\alpha}(2\gamma_V - 1) - 8\delta\bar{\alpha}^3\psi\gamma_V(1 - \gamma_V) > 0 \quad (37)$$

These reduce to

$$x^L < \frac{-\bar{\alpha}(2\gamma_V - 1) - (1 + \delta)8\bar{\alpha}^3\psi\gamma_V(1 - \gamma_V)}{8\bar{\alpha}^2\psi(\gamma_V - \gamma_L) - 1} \quad (38)$$

$$\gamma_L < \gamma_V - \frac{1}{8\bar{\alpha}^2\psi} \quad (39)$$

$$\bar{\alpha}^2 > \frac{1}{8\psi\gamma_V} \quad (40)$$

$$x^R > \frac{\bar{\alpha}(2\gamma_V - 1) + (1 + \delta)8\bar{\alpha}^3\psi\gamma_V(1 - \gamma_V)}{8\bar{\alpha}^2\psi(\gamma_R - \gamma_V) + 1} \quad (41)$$

$$\gamma_R > \gamma_V - \frac{1}{8\bar{\alpha}^2\psi} \quad (42)$$

□

Proposition 6: Any pair of platforms satisfying:

1. $E[U_V(x_1^{L*})] = E[U_V(x_1^{R*})]$

2. $\underline{x} \leq x_1^{L*} \leq x_V^g \leq x_1^{R*} \leq \min \in \{\bar{x}, x_L^{pos}, \tilde{x}\}$, where \tilde{x} is the maximum of R 's expected utility in the range $[0, x']$

can be sustained in a a gambling equilibrium.

Proof. The proof proceeds as for Proposition 2 and is therefore omitted. □