The London School of Economics and Political Science

Playing With Fire or Playing It Safe? Formal Models of Gambling in Elections

Federica Izzo

A thesis submitted to the Department of Government of the London School of Economics and Political Science for the degree of Doctor in Philosophy

London, April 2019

Declaration

I certify that the thesis I have presented for examination for the MPhil/PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it). The copyright of this thesis rests with the author. Quotation from it is permitted, provided that full acknowledgement is made. This thesis may not be reproduced without my prior written consent. I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party. I declare that my thesis consists of 39.000 words.

Acknowledgements

First of all, I thank my supervisors Stephane Wolton and Torun Dewan. There are simply no words to describe how indebted I am to Stephane for his patience, commitment and constant support. No one has dedicated more time to my work, and no one has pushed me and motivated me more than he has done in the past five years. He has been a mentor and a friend, and without him my PhD experience would have been infinitely less fun. I am forever grateful to Torun for 'showing me the righteous path' so many years ago, and for believing that I was a gamble worth taking. Quite simply, his support form that moment onwards has changed my life.

I am indebted to Gilat Levy and Ethan Bueno de Mesquita for their invaluable support and their help in developing my research. I am also grateful to Ethan for sponsoring me as a visiting student at Harris. I enjoyed my time in Chicago immensely. I also thank Catherine Hafer and Dimitri Landa for hosting me at NYU, and for making my last visit there so much better!

I also want to thank the friends that made this otherwise lonely endeavour such a fun journey: Clement(e), Selina, Arduino, Greg, Alberto, Jan, Tinghua, Ali, Kostas, Janne, Giovanna, Jesse, Kattalina, Korhan and Amanda. A special thank you goes to Nicola Mastrorocco, without whom I would have probably gone insane many years ago.

Ringrazio immensamente la mia famiglia. Mio padre, perchè ha sempre creduto che non ci fosse nessun ostacolo che io non fossi in grado di superare. E' solo grazie a lui che ho avuto il coraggio di affrontare le sfide incontrate finora, grandi o piccole che fossero. Mia madre, perchè riesce a capirmi come nessun altro al mondo. Quando ho bisogno di un sorriso, so sempre che lo troverò da lei. E mia sorella, per il suo affetto incondizionato. La amo immensamente, e non c'è nulla al mondo che mi dia più gioia che vederla felice.

In fine, ringrazio mio marito. E’ a lui che questa tesi è dedicata. E’ l’amore della mia vita, e senza averlo al mio fianco questo dottorato probabilmente non l’avrei mai iniziato. Il suo amore, il suo supporto incondizionato, la sua fede cieca nelle mie capacità sono la mia più grande fonte di forza. Senza nessuna esitazione, ha messo i miei desideri e le mie ambizioni davanti alle sue. Per questo, gli sarò per sempre grata.

Abstract

When voters face uncertainty over their optimal choice, the outcome of today's policy making influences their future beliefs. This generates incentives for politicians to engage in information control. In this dissertation, I formally analyze how these incentives influence the conflict both within and between political parties.

My first paper begins with the observation that political leaders are often publicly attacked by their own ideological allies. Yet, evidence indicates that this form dissent is electorally costly, thus harming both the leader and his allies. Why, then, does it emerge in the first place? I address this puzzle within a model in which voters face uncertainty about their ideal policy and learn via experience. In particular, I propose a new framework to think about policy experimentation, whereby the amount of learning depends on the location of the implemented policy on the left-right spectrum. I show that, within this setting, dissent emerges precisely because it is electorally costly. By hurting the incumbent's electoral chances, dissent alters his incentives to adopt more or less informative (and therefore extreme) policies. This creates a potential trade-off for the incumbent's allies which, under some conditions, is resolved in favour of dissent.

This policy experimentation framework can be adapted to study several phenomena. In my second paper, I apply it to investigate whether ideological parties may have strategic incentives to lose elections. Parties often take extreme positions even if this means losing for sure. Extant explanations rely on expressive motivations. I instead show that a party whose ideological stance is ex-ante unpopular faces a trade-off between winning the upcoming election and changing voters' future preferences. Under some conditions, the party chooses to lose today to win big tomorrow.

Finally, my third paper focuses on electoral selection, addressing a crucial question: do the right candidates for office choose to run at the right time? In the model, voters learn about office holders' competence by observing governance outcomes. Because competence matters most in times of crisis, this is also when outcomes reveal most information. I show that electoral accountability has the perverse consequence of discouraging good candidates from running in times of crisis, that is, precisely when voters would need them the most.

Contents

Page No.

Playing with Fire or Playing it Safe: a Brief Introduction **8**

With Friends Like These, Who Needs Enemies? **16**

Introduction . . . 17
Related Literature . . . 21
The Model Set-Up . . . 24
Equilibrium Analysis . . . 28
Learning and Experimentation . . . 29
The Incumbent . . . 32
Dissent by an Exreme Ally . . . 35
Welfare Analysis . . . 41
Micro Founding the Electoral Cost of Dissent . . . 43
Extension: What if the Ally Has Bargaining Power . . . 45
Empirical Implications and Falsifiability . . . 48
Conclusion . . . 52
References . . . 55
Appendix A: Proofs . . . 59
Appendix B: Dissent by a Moderate Ally . . . 84

Ideology for the Future **88**

Introduction . . . 89
Literature Review . . . 93
The Model . . . 95
Heterogeneous Priors and Beliefs as Ideology . . . 97
Analysis: Learning . . . 99
The Voter . . . 101
The Parties . . . 102
The Parties' Utility . . . 103
Gambling on the Future . . . 105
Parties' Beliefs and Ideology . . . 112
A Look at a Forward Looking Voter . . . 115
References . . . 117
Appendix . . . 123

Do We Get the Best Candidates
When We Need Them the Most? **135**

Introduction . . . 136
Literature Review . . . 139
The Baseline Model . . . 140
Analysis . . . 144

Discussion and Robustness . . . 149
The Electoral Effect of Incumbency . . . 152
Isolating the Information Channel . . . 155
Conclusion . . . 160
References . . . 161
Appendix A: Proofs . . . 163
Appendix B: Robustness . . . 170

List of Figures

Page No.

With Friends Like These, Who Needs Enemies?

Figure 1: Voter's Payoff . . . 30
Figure 2: Effect of Dissent on Equilibrium Policy . . . 38
Figure 3: Probability of Winning . . . 49
Figure 4: Probability of Winning Observable . . . 50
Figure 5: Probability of Winning Observable, $x^A > \underline{x^A}$ . . . 51

Ideology for the Future

Figure 1: Voter's Payoff . . . 100
Figure 2: Party L's first-period utility . . . 104
Figure 3: Party L's expected overall utility . . . 105
Figure 4: Players' utility . . . 108
Figure 5: Forward looking voter's expected overall utility . . . 117

Playing with Fire or Playing it Safe, a Brief Introduction

Free and fair elections are the defining feature of a democratic system. Ideally, elections should in fact allow 'the people' to both select the best candidate for office, and incentivise him or her to undertake their preferred course of action. However, the political world is complicated: there is often substantial uncertainty over what is the optimal choice for the voters. This uncertainty may refer to which policies best suit the voters' interests, to which candidate is most capable of advancing them, or both. Indeed, it is often hard to anticipate the exact consequences of each policy choice. Similarly, it is difficult to tell ex-ante whether a candidate possesses the right qualities to effectively deal with a situation of crisis.

There are two ways in which voters may potentially solve this uncertainty. First, they may learn via information transmission. Politicians often possess information that is unavailable to the voters. Then, by observing what politicians say or do, voters update their beliefs over which policy or candidate is the 'correct' one. However, it is not always the case that politicians 'know better'. Politicians and voters may often have access to the same information on how the world works. Even when they have different views of the world, these may be the result of different ideologies rather than information asymmetries. Then, the only way that the voters can learn is via experience. When facing uncertainty over which platform would produce the best outcome, voters turn to the consequences of today's policy making to refine their expectations, accordingly update their preferences and consequently modify their electoral choices. Similarly, when the uncertainty refers to politicians' true ability, the incumbent's performance in dealing with complicated tasks allows voters to review their assessment of his competence relative to the other candidates'.

A large strand of the formal literature in political science and political economy focuses on the consequences of voters' uncertainty under the assumption that learning occurs via information transmission. In contrast (with few exceptions, that are discussed in more detail in the chapters below), very little attention is devoted to understanding how this uncertainty influences politicians' strategic behaviour (and, in turn, electoral and policy outcomes) when voters learn via experience. This question is precisely what this thesis is concerned with.

This thesis builds on a simple intuition: when voters face electorally relevant uncertainty and learn via experience, the amount of information they observe is a function of the politicians' strategic behaviour. As such, politicians may face a trade-off between taking their preferred action today and generating the optimal amount of information to maximize their future electoral chances. They may have incentives to *gamble* by taking an action that allows voters to learn, or may instead *fear* that information would hurt them and thus prefer to play it safe. In the **three papers** of this thesis I present formal models to analyse how these incentives influence the **ideological conflict** both **within** and **between political parties** (respectively, in *'With Friends Like These, Who Needs Enemies?'* and *'Ideology for the Future'*), as well as the **quality of political candidates** (in *'Do We Get the Best Candidates When We Need Them the Most?*). In the remainder of this introduction I provide a brief summary of the three papers, thus highlighting both the methodological and substantive contributions of the thesis.

With Friends Like These, Who Needs Enemies?

As recent events have highlighted, deep divisions exist within political parties in the US. Similarly, open conflict between competing factions characterizes many European parties. Italy's Democratic Party, the French Socialists and the British Conservatives offer obvious examples. Yet, the importance of factionalism and intra-party dynamics is often overlooked in the formal theory literature. My first paper *'With Friends Like*

These, Who Needs Enemies?' contributes to this literature by analysing the *causes* and *consequences* of the emergence of *open conflicts within political parties*. While the few extant works focus on intra-party cohesion in legislatures, my paper begins with the observation that intra-party dissent often takes the form of *public attacks* against the party leader by his own ideological allies (such as a minority faction within the party, or an ideologically aligned media outlet). Yet, evidence indicates that this form of dissent damages the party's electoral chances. As such, it harms both the leader and the dissenters themselves. Why, then, does it emerge in the first place?

I argue that this form of public dissent occurs precisely because it is electorally costly, in order to induce a policy response. Dissent hurts the leader's chances of winning re-election. This, in turn, changes his incentives to take policy gambles. As such, when gambles take the form of more extreme policies (as they will do *endogenously* in my framework), the allies face a potential trade-off between maximizing the probability that the leader wins the upcoming elections and inducing him to adopt a policy more in line with their own ideological preferences. If the gain from changing today's equilibrium policy is sufficiently large, public dissent emerges in equilibrium.

In the paper I micro-found this argument within the context of a principal-agent model. The incumbent's ideological allies (that come from the same side of the political spectrum but do not share exactly the same policy preferences) must decide whether to attack him, thus damaging the party's electoral chances, or stay quiet. The incumbent then chooses which policy to implement. After the voter observes the outcome of the implemented policy, she chooses whether to retain the incumbent or replace him with this challenger. The key innovation of the paper is to consider a setting in which the voter faces uncertainty about her own policy preferences, and learns via experience. In order to analyse a world with these features, I propose a new framework to think about policy experimentation, whereby the amount of learning depends on the location of the implemented policy along the left-right spectrum. Voters learn about their own ideological preferences by observing how much they like the outcome of today's policy. The presence of a random shock complicates their inference problem. Within this framework, I show, voters learn more about their ideal policy when extreme platforms are implemented. The extent to which voters' preferences change over time then depends *endogenously* on the incumbent's policy choice.

As a consequence, the office holder has incentives to engage in information control. Leading incumbents (that have high chances of being re-elected) experience *fear of failure*: they have incentives to adopt moderate policies that prevent information generation. In contrast, trailing ones (that have a low re-election probability) have incentives to *gamble for resurrection*: they want to engage in extreme policies that allow voters to learn, in hopes of improving their electoral prospects. Within this setting, dissent emerges precisely because it is electorally costly. By dissenting, the incumbent's allies hurt his chances of winning the upcoming election. This, in turn, alters his incentives to implement more or less informative (i.e. extreme) policies. As such, the allies face a potential trade-off between maximizing the incumbent's re-election chances and inducing him to implement a platform more in line with their own ideological preferences. If the gain from changing today's equilibrium policy is sufficiently large, optimally balancing the trade-off involves active dissent that damages the party electorally.

The paper then concludes by analysing the normative and empirical implications of the results. First, I show that the presence of the incumbent's extreme ally can be welfare improving for the voter. By dissenting the ally can in fact induce an optimal level of policy experimentation by the office holder, thereby mitigating an inefficiency that has its roots in electoral accountability. Finally, the results also have important implications for empirical research, showing that existing estimates of the electoral rewards of party unity are inevitably biased.

Ideology for the Future

In my second paper, I move from intra-party to inter-party conflict. My starting point is the observation that political parties sometimes take extreme positions, even if this means losing for sure. Existing explanations rely on the assumption that their members have expressive rather than strategic motivations, and care about ideological purity. Corbyn's election as leader of the UK Labour party has been interpreted in this light.

Similarly, a desire for ideological purity is assumed to lie at the roots of the 'Sanders phenomenon' in the US. In *'Ideology For the Future'*, I ask whether ideological parties may instead have *strategic* incentives to lose election, even absent any concern for purity. I argue that a party whose ideological stance is unpopular with the electorate faces a trade off, 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 therefore gambles on the future: chooses to lose today to change voters' views and win big tomorrow.

I micro-found this argument by presenting a model of repeated spatial elections with two time periods, in which I embed the policy experimentation framework developed in my first paper. A key feature of the model is that the players have different priors over the location of the voter's true ideal policy (i.e., the state of the world), but agree to disagree. In particular, I think about the parties' beliefs as a second dimension of their ideology. Thus, parties truly believe that their own ideological preferences are in line with the state of the world. Crucially, the voter's ex-ante preferences (i.e., her prior beliefs over the location of her ideal policy) are common knowledge. This implies that, given any pair of platforms, parties face no uncertainty over the outcome of the upcoming election. However, uncertainty – and, given heterogeneous priors, disagreement – exist over what the voter will learn upon observing the first period policy outcome.

Given these assumptions, the second period election is equivalent to a one-shot Downsian game: the parties always converge on the voter's preferred policy. Not so much in the first period. 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 unpopular 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 voters 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 a 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, it shows that 'extremism' in both preferences and beliefs is necessary: ideological beliefs are therefore an essential part of the story.

Do We Get the Best Candidates When We Need Them the Most?

While the first two chapters focus on the ideological conflict that emerges both between and within parties, my third paper considers a world in which political candidates differ only in their expected ability. As several empirical paper highlight, the competence of political leaders has a crucial impact on a country's performance. As such, it is important to understand under which conditions high-quality politicians are willing to run for office in the first place. One question is particularly relevant to evaluate the effectiveness of democratic elections in improving voters' welfare: do the right candidates self-select at the right time? More specifically, are the most competent politicians willing to run for office during times of crisis, when competence matters the most? This is the question that I address in 'Do We Get the Best Candidates When We Need Them the Most?'.

The formal literature has so far placed little emphasis on the topic of political self-selection. Most extant models of elections in fact take the pool of candidates as exogenous, focusing instead on voters' ability to identify good politicians to be (re)elected and bad ones to be thrown out. A small recent literature allows for endogenous candidate entry, thereby analysing the equilibrium supply of good politicians. However, these works typically consider a static setting, highlighting the difficulty of attracting competent politicians if office rents are too low compared to private market salaries.

In this paper, I adopt a very different perspective. I consider a world in which potential candidates are career politicians, for whom office is always more valuable than the outside option. As such, entering the race is always the statically optimal choice for all potential candidates, irrespective of their expected ability and the conditions in the country. I show that this does not always hold true when we take into account their dynamic incentives. Under some conditions, 'good' candidates are not willing to run for office during times of crisis.

This result emerges within the context of a model of repeated elections in which potential candidates differ in the probability of being competent. The crucial element of the set-up is that in each period the country either experiences a crisis, or a period of 'business as usual'. A crisis (economic or otherwise) is an *exogenous* shock that has two key features: it amplifies the impact of the office holder's competence and, at the same time, the informativeness of his performance. In other words, precisely because competence matters the most during times of crisis, this is also when the governance outcomes reveals most information about the office holder's ability.

Within this framework, all potential candidates are willing to run for office during normal times. Not so much during periods of crisis. In fact, the politician who is most likely to solve the crisis also has the most to lose from failing. This politician in fact has a valuable electoral advantage. As such, he would want to prevent the voter from learning about his true ability, since information can only hurt his future electoral chances. Unfortunately, it is precisely during times of crisis (that is, when competence matters the most) that governance outcomes reveal most information about the office holder's type. Then, the politician who is most likely to be competent experiences *fear of failure* and chooses to stay out of the race in order to preserve his electoral capital for the future. In contrast, the 'worst' (in expectation) potential candidate has nothing to lose. He is therefore always willing to take the *gamble*, and run for office during challenging times. Thus, the voters gets the *wrong candidates* at the *wrong time*. Crucially, this adverse selection does not arise due to weak electoral incentives, as it is the case the extant literature. Quite the opposite, it emerges precisely as a perverse consequence of accountability.

With Friends Like These, Who Needs Enemies?

Abstract

Why are political leaders often attacked by their ideological allies? The paper addresses this puzzle by presenting a model in which the conflict between the incumbent and his allies is ideological, dissent is electorally costly, and voters are learning about their own policy preferences over time. Here, by dissenting against the incumbent (and thereby harming the party in the upcoming election), the allies can change his incentives to choose more or less extreme policies, which affects the amount of voter learning. This induces a trade-off between winning the current election and inducing the party leadership to pursue the allies' all-things-considered more-preferred policy. Optimally balancing this trade-off sometimes involves active dissent that damages the party in the short-run. In equilibrium dissent arises precisely because it is electorally costly.

Introduction

'Renzi is not apt for his role, he does not have the stature of a leader' (Cuperlo 2016). 'He says a lot of things, that do not always coincide with the truth' (Bersani 2015). These are public statements made by prominent Italian politicians about Matteo Renzi, former prime minister and leader of the Democratic Party (PD), the biggest Left-wing party in the country. And these are not isolated examples: Renzi was often publicly accused of being a liar, incompetent and even 'worse than the devil' (D'Alema 2016). Quite surprisingly, the authors of these public attacks were not members of the opposition. Renzi's worst critics were in fact all members of his own party: the leaders of the so called "Minoranza Dem", the extreme minority faction within the PD.

Similar phenomena have emerged in other European countries as well as in the US. In the UK, the Labour Party is currently undergoing a 'civil war' (Jones 2016). The members of the party's minority often openly denigrate the leader Jeremy Corbyn whom, they argue, 'literally has no idea (...) how to conduct himself as a leader' (Mandelson 2017). In France, a group of rebel Socialist MPs (the Frondeurs) regularly manifested their dissent against President and party leader Francois Holland. Similarly, divisions within the US Republican party are apparent. Prominent Republican lawmakers have publicly attacked President Trump, arguing that 'he shows a growing inability, and even unwillingness, to separate truth from lies' (McCain 2017).

Interestingly, this phenomenon is not solely an intra-party issue. Media outlets often denigrate political leaders with whom they are ideologically aligned. The right-leaning Evening Standard has openly attacked UK Conservative prime minister Theresa May, depicting her cabinet as 'stale' and 'enfeebled' (Urwin 2017). Similarly The Guardian, historically left-leaning, has described Labour leader Corbyn as 'dismal, lifeless, spineless' (Toynbee 2016).

These examples show that political leaders are often publicly attacked and denigrated by their own ideological allies. Yet, evidence indicates that this form of dissent typically damages a party's electoral chances, since voters dislike parties that appear divided (Greene and Haber 2016; YouGov 2016; Kam 2009; Groeling 2010). As such, public dissent hurts both the leader and the dissenters themselves. This raises the question: why would a leader's ideological allies choose to publicly attack him despite this being electorally harmful?

In this paper, I argue that this form of public dissent emerges precisely because it is electorally costly, in order to induce a policy response. Dissent hurts the leader's chances of winning re-election. This, in turn, changes his incentives to take policy gambles. As such, when gambles take the form of more extreme policies (as they will do *endogenously* in my framework), the allies face a potential trade-off between maximizing the probability that the leader wins the upcoming elections and inducing him to adopt a policy more in line with their own ideological preferences. If the gain from changing today's equilibrium policy is sufficiently large, public dissent emerges in equilibrium.

Focusing on dissent against an incumbent, I micro-found this argument by presenting a model with four key ingredients. First, the incumbent and his allies come from the same side of the political spectrum, but do not have exactly the same policy preferences. The allies can represent a minority faction within the party, a media outlet, an external donor or even a special interest group: any actor whose policy preferences are closer to the incumbent's than the challenger's. Second, dissent is electorally costly: it generates a negative valence shock that potentially damages the party's electoral prospect. Dissent can entail publicly criticizing the party leader, dispraising his policy choices, revealing a scandal or even 'mechanically' reducing his electoral chances (for example, a donor may choose to reduce its electoral contributions). Third, the model assumes that voters face uncertainty about their ideal policy. For example, voters may not know which policy is most likely to produce their desired outcome. Finally, a crucial feature of the model is that voters can reduce their uncertainty by learning through experience. In particular, I propose a new framework to think about policy experimentation. Voters learn about the optimal decision for the future by observing how much they like the outcome of today's policy. The presence of a random shock complicates their inference problem. Within this framework, I show, the amount of voter learning depends on the location of the implemented policy along the left-right spectrum. The more extreme the policy is, the more the voters learn about their ideal platform. Suppose that an extreme policy is implemented. If a voter obtains a high (low) payoff from the resulting outcome, the policy is likely (unlikely) to be in line with her true preferences. Conversely, because of the presence of the random shock, the outcome of a moderate policy is much less informative.

In this setting, the incumbent has incentives to engage in information control. His equilibrium policy choice maximizes the trade-off between implementing his bliss point today and generating the optimal amount of information in order to be re-elected tomorrow. This, I show, is a function of the incumbent's ex-ante electoral strength. A leading incumbent, who is going to be re-elected even if the voters receive no new information, has incentives to implement moderate platforms that prevent information generation. In contrast, a trailing one will want to engage in extreme policies that increase the amount of voter learning, in hopes of improving his electoral prospects. Finally, an incumbent who can never be re-elected (irrespective of what the voters learn) is indifferent with respect to the amount of information that is generated, and will simply follow his ideological preferences.

Within this framework, dissent may allow the allies to solve the ideological conflict with the incumbent. By dissenting, the allies generate a negative valence shock against the incumbent, thereby reducing his ex-ante electoral strength. This, in turn, creates incentives to implement more or less informative (i.e. extreme) policies. As such, dissent changes the incumbent's equilibrium policy choice, while also harming the party electorally. This generates a potential trade-off for the incumbent's allies, between ensuring that their preferred party wins the upcoming election and inducing the incumbent to implement a policy more in line with their own ideological preferences. Optimally balancing this trade-off sometimes involves active dissent that damages the party in the short run. Thus, dissent emerges precisely because it is electorally costly, and it produces unity of interests between the incumbent and his allies even if no player actually cares about unity per se. Surprisingly, the analysis reveals that improving the incumbent's electoral prospects or reducing his ideological conflict with the allies may make dissent more likely to emerge.

Further, the results highlight that the presence of an extreme ally to the incumbent party may be welfare improving for the voters. In the model, voters benefit from informative policies being implemented as this increases the probability of making the correct electoral decision in the future. However, under some conditions, electoral accountability has the perverse consequence of inducing lower levels of policy experimentation relative to both the incumbent's ideological preferences and the voter's optimum. The incumbent's extreme ally may mitigate such inefficiency. By dissenting, the ally can create incentives for the incumbent to implement extreme policies that allow the voters to learn. If the value of acquiring new information is sufficiently large, this strictly increases the voters' welfare.

The results of the model also have an important implication for empirical research on the topic. Existing estimates of the electoral rewards of party unity, that are obtained by comparing treated and control units (i.e. parties that do and do not experience dissent), are inevitably biased. In addition, it is hard to know ex-ante what the direction of the bias will be. However, this does not imply that the model is not falsifiable. Indeed, the theory suggests where else to look in order to empirically investigate the electoral consequences of dissent. The model generates testable comparative statics regarding parties' electoral performance conditional on experiencing dissent. Focusing on this restricted sample, and thereby avoiding the problem of selection bias described above, researchers can empirically investigate the conditions under which dissent is expected to hurt parties the most. The theory predicts that parties' performance conditional on dissent should be positively correlated with variables such as the level of education, news media consumption and political engagement in the electorate. Finally, I discuss how the model's comparative static predictions may allow us to distinguish it from other possible explanations for the emergence of dissent.

Related Literature

This paper relates first and foremost to the literature on intra-party politics. In the formal literature, the interaction between different factions is typically analysed as a bargaining game. Mutlu-Eren (2015) considers how the threat of a split influences the party's behaviour in the legislature. Similarly, Hortala-Vallve and Mueller (2015) consider a model in which the threat of defection by the minority can induce the party leadership to democratize the candidates' selection process. In these papers, the threat is credible when the faction is sufficiently likely to win the upcoming election if running alone after a split. Turning to the empirical literature, we find similar references to the competing factions bargaining over a prize. In Parties and Party Systems, Sartori (1976) describes factions as blackmailing the leadership, and seeking side payments. Belloni (1976) and Boucek (2009) express similar ideas. More recently Budge et al. (2010) explain parties' policy shifts away from the center as a result of the minority faction vetoing a moderation.

Yet this approach has some issues when we consider dissent as *public attacks* against the leader, rather than dissent as formal defection. In a bargaining game dissent would be used as a threat, to be executed *after* the incumbent has made his policy choice. However, at this point dissent has no effect but to reduce the probability that the party wins the upcoming election. This strictly decreases both the incumbent's and the ally's expected payoff. As such, the threat can never be credible and we should never observe dissent in equilibrium. Further, even beyond the issue of credibility, in a bargaining game the materialization of the threat typically lies off the equilibrium path. Hence, this is arguably not an appropriate framework to understand why political parties so often experience open dissent.

A second strand of literature considers politicians' incentives to pander to their own individual constituencies (e.g., Carey and Shugart 1995 and Buisseret and Prato 2018). Politicians may face a trade-off between the national party's electoral fortunes and their own success. In particular, this trade-off may emerge if a politician's local constituency is opposed to the national party line. Within this framework, dissent may serve the purpose of signaling the politician's misalignment with the leadership (and thus alignment with the constituency's preferences).

This argument is certainly intuitive (and potentially applies to formal defections as well as public attacks against the leadership). However, both anecdotal and systematic evidence seem to suggest that it does not fully capture the rationale behind this phenomenon. In particular, according to this framework we should expect the individual dissenters to be in a relatively weak electoral position (i.e., not to come from safe constituencies). However, Proksch and Slapin (2015) analyze data from the UK and Germany and show that, if anything, the opposite holds. Public attacks against the leadership are (weakly) more likely to come from members of parliament elected with a larger margin. Further, the authors exploit the features of the German mixed-member proportional electoral system, and show that members elected with a party-list vote are as likely to dissent as those elected with a constituency vote. Similarly hard to reconcile with the 'pandering' argument is the emergence of dissent under a closed-list system such as the Italian one, where the party leader controls the list composition and as such the dissenter's electoral fate. This calls for a theory that allows us to make sense of this phenomenon even when the party members' individual electoral motives do not provide incentives (or worse, provide disincentives) to attack the leadership.

Thus, this paper presents a substantially different type of model, in which dissent precedes rather than following the party leader's strategic choice (in contrast with a bargaining set-up), and emerges in order to induce a policy response (in contrast with the pandering set-up).

The core ingredient of the model is the voters' uncertainty over their optimal choice. Given the symmetric lack of information, such uncertainty may only be resolved via experience. This connects the paper with the research on learning and experimentation. The key intuition therein is that, when deciding which policy to implement today, politicians consider how the outcome will influence the voters' future beliefs. Most extant works assume that the voters must learn about the incumbent's type, i.e. his ability or competence. The incumbent chooses between a safe and a risky policy, with a success on the latter being conditional on the politician being a 'good type' (see for example Dewan and Hortala-Vallve 2016; Majumdar and Mukand 2004). Under the assumption of symmetric uncertainty, a risky policy is always a gamble. This paper differentiates itself from the extant literature by considering a setting in which the incumbent's incentives to gamble arise endogenously from his allies' strategic behaviour. Further, the voters must learn about their own policy preferences (i.e. the state of the world), and not about the office holder's competence.

In this perspective, the paper is closely related to recent work by Callander (2011). The author considers a world in which players face uncertainty about how policies map into outcomes: they know the slope of the mapping function (representing the state of the world), but try to fine-tune their predictions by learning about the exact realization of the variance. The nature of the uncertainty is reversed in this paper: the voters must learn the fundamental underlying state of the world. This generates the result that extreme policies, rather than small incremental changes as in Callander, produce more information. As such, this paper provides a new framework to think about policy experimentation. Additionally, Callander focuses on the statically optimal choice for a decision maker. He thus chooses to abstract from dynamic considerations, by assuming either myopic players (Callander 2011) or exogenous retention probabilities (Callander and Hummel 2014). In contrast, the focus of this paper is precisely on the incumbent's dynamic incentives to control information, and on how these impact his policy choices and the conflict with his ideological allies.

Finally, the paper relates to the literature on Bayesian Persuasion, originated from the work of Austen-Smith (1998) and Kamenica and Gentzkow (2011). In my model, as in the Bayesian Persuasion framework, the incumbent can engage in information control by manipulating the receiver's posterior distribution. In the Bayesian Persuasion framework the mechanism through which this occurs is somewhat black-boxed. The assumption is that the persuader can choose any desired signal (i.e., realization space and conditional probability distributions) by designing a 'test' or a 'policy experiment'. The results of the test do not directly influence the players' payoffs, therefore the persuader must be able to credibly commit to truthfully revealing them to the public. The key innovation of my paper is therefore to explicitly model *how* the incumbent can engage in information control, by looking at the impact that the *implemented policy* has on voter learning.

The Model Set-Up

Dissent is analysed within the framework of a principal-agent model, under the assumption that the voters face uncertainty over their ideal policy (the state of the world) and learn by experience. I focus on dissent within the incumbent party. The players are therefore the incumbent (I), his ideological ally (A), a challenger (C), and a representative voter (V). The incumbent's ally can represent a minority faction within the party, media outlets ideological close to the incumbent, donors or even interest groups: any actor whose ideological preferences are closer to the incumbent's than the challenger's.

At the beginning of the game, the incumbent's ally chooses whether to dissent against him. The choice is binary: $D \in \{0, 1\}$. Dissent may entail publicly criticizing the incumbent's personality, or manifesting a disagreement with the party line. After observing his ally's choice, the incumbent implements a policy x_1 along the real line. The voter chooses whether to retain the incumbent or replace him with his challenger. The second-period office-holder implements a new policy x_2 (under the assumption of no credible commitment).

The voter faces uncertainty over the exact value of her ideal policy x^v.[1] One way to interpret this assumption is that the voter does not know which policy is most likely to produce her preferred outcome. Thus, her uncertainty refers to the slope of the function mapping policies into outcomes. An alternative interpretation would consider a voter that knows how different policies map into outcomes, but can not perfectly anticipate how these outcomes will impact her payoffs. The voter's ideal policy can take one of two values: $x^v \in \{\underline{\alpha}, \bar{\alpha}\}$. For simplicity (but without loss of generality) I assume $\underline{\alpha} = -\bar{\alpha} < 0$. The qualitative results survive if $\underline{\alpha}$ and $\bar{\alpha}$ have the same sign, that is if the voter knows whether her ideal policy is a left-wing one or a right-wing one, but faces uncertainty over its exact location.

The model features no asymmetry of information: no player knows the true value of x^v, and all players assign the same prior probability γ to the voter's ideal policy being a right-wing one ($\gamma = \text{prob}(x^v = \bar{\alpha})$). Given this symmetric uncertainty, learning only happens via experience. The voter observes how much she liked (or disliked) the first period policy, and updates on the true value of x^v by using Bayes rule.[2] Formally, the

[1] While the model only considers a representative voter, the results do not require all voters to face such uncertainty. Indeed, some voters may be ideological and have well-defined policy preferences. The results presented below go through as long as the 'uncertain' voters are pivotal in determining the electoral outcome.

[2] Whether the incumbent, his allies, and the challenger also observe the voter's payoff realization

voter's payoff realization is a noisy signal of the state of the world:

$$U_t^v = -(x^v - x_t)^2 + \epsilon_t - \mathbb{I}\delta \tag{1}$$
$$\epsilon_t \sim U[-\frac{1}{2\psi}, \frac{1}{2\psi}]$$

As I will discuss in more details below, the assumption that the random shock ϵ is uniformly distributed is not necessary for the results. The parameter δ captures the observation that, everything else being equal, voters dislike parties that appear divided: if the incumbent experiences dissent in the first period, the voter's expected utility from re-electing him is reduced by δ ($\mathbb{I} = 1$ if $D = 1$ and the incumbent is re-elected and $\mathbb{I} = 0$ otherwise). In other words, I assume that dissent generates an endogenous valence shock against the party. In order to simplify the analysis and presentation of the results, I leave the cost of dissent black-boxed. I will discuss possible micro-foundations of this assumption in a separate section.

Finally, I A and C are policy motivated, and their bliss points are common knowledge:[3]

$$U_t^i = -(x^i - x_t)^2 \quad \forall i \in \{I, A, C\} \tag{2}$$

Without loss of generality, I will consider a right-wing incumbent and a left-wing challenger: $x^C \leq 0 \leq x^I$. For simplicity, I also assume that the candidates' bliss points are symmetric around 0: $x^I = -x^C \geq 0$. The incumbent and his ally come from the same side of the ideological spectrum (i.e. are both right-wing), but do not have

is inconsequential for the equilibrium results.

[3]For purposes of presentation (and to focus on learning via experience) the model assumes away any asymmetry of information. However, the key insights of the paper would (under some conditions) survive if we allow the incumbent to have private information about his own preferences or the voter's bliss point.

exactly the same bliss point. However, the ally's preferences are always closer to the incumbent's than to the challenger's:

$$|x^A - x^I| < |x^A - x^C| \tag{3}$$

In the main body of the paper I will focus on the case of an extreme ally ($x^A > x^I$).

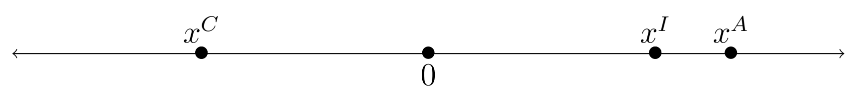

In the Appendix, I show that within this framework dissent can emerge even when the ally is more moderate than the incumbent, and identify the conditions under which this occurs in equilibrium.

Timing

1. Nature determines the value of $x^v \in \{\underline{\alpha}, \bar{\alpha}\}$
2. The Incumbent's Ally chooses whether to dissent against him: $D \in \{0, 1\}$
3. The Incumbent implements a policy $x_1 \in \mathbb{R}$
4. The Voter's first-period payoffs realize
5. The Voter chooses whether to re-elect the Incumbent or replace him with the Challenger
6. The second-period office holder implements policy $x_2 \in \mathbb{R}$ (no credible commitment)
7. Second-period payoffs realize and game ends

The equilibrium concept is Perfect Bayesian Equilibrium. In order to avoid trivial results, I assume that when indifferent the incumbent's ally chooses not to dissent. This is formally equivalent to assuming an infinitely small material cost of dissenting.

In order to isolate the impact of ideological disagreements, I do not include office rents in the players' utility function. Whenever the incumbent and his ally do not attach the same value to winning office per se, office rents would in fact represent a second source of conflict. Suppose for example that the ally represents a minority faction within the party. Should the party win the upcoming election, the incumbent (i.e. the leader of the majority faction) would arguably grab a larger share of the office rents relative to his ally. This potentially translates into different risk appetite in policy making, thereby increasing the conflict of preferences between the incumbent and his allies. Hence, as long as the value of office is not too large, including office payoffs would make dissent even easier to sustain in equilibrium.

Equilibrium Analysis

As usual, we proceed by backwards induction, starting from the second period's office holder's choice. Politicians have no credible commitment ability. As such, given the absence of re-elections incentives, the second period office-holder will always implement his preferred platform. The voter therefore faces a selection problem. Her electoral choice will then be determined by the (posterior) beliefs that her own ideal policy is aligned with the incumbent's preferred platform, as well as by the presence or absence of dissent within the incumbent party. Specifically, in any PBE of the game, the voter re-elects the right wing incumbent if and only if the posterior probability of being ideologically aligned with him ($\mu = prob(x^v = \bar{\alpha})$) is sufficiently high:

$$\mu > \frac{\mathbb{I}\delta + 4\bar{\alpha}x^I}{8\bar{\alpha}x^I} \tag{4}$$

The indifference breaking assumption is without loss of generality. Notice that, absent dissent, the incumbent is always re-elected as long as $\mu > \frac{1}{2}$. When the incumbent experiences dissent, the higher the cost δ, the higher the voter's posterior needs to be to guarantee re-election

Learning and Experimentation

Moving one step backwards, consider the voter's inference problem. The voter observes how much she liked or disliked the first period policy, and updates her beliefs on her ideal policy by using Bayes' rule. The analysis reveals a crucial feature of the learning process: the amount of information obtained by the voter depends on the location of the policy implemented in the first period. Specifically, the voter learns more from more extreme policies. As the implemented policy becomes more extreme, the distance in the expected outcomes 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, such policy is likely (unlikely) to be in line with her true ideology. However, given the presence of the random shock, the outcome of a moderate policy is much less informative. This feature emerges in a very stark form in a world in which the shock is drawn from a uniform distribution.

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

(i) Her posterior μ takes one of three values: $\mu \in \{0, \gamma, 1\}$;

(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 $\mu \neq \gamma$ with probability 1.

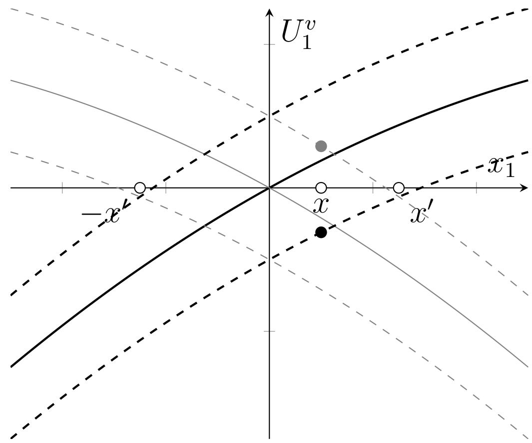

Figure 1: Voter's payoff realization as a function of first-period policy. The thick increasing (thin decreasing) curves represent the case in which $x_V = \overline{\alpha}$ ($x_V = \underline{\alpha}$). The solid curves represent the voter's expected payoff $E[U_1^v]$, while the dashed ones represent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

Lemma 1 tells us that the voter either learns everything or nothing. Further, the probability that the voter discovers her true preferences increases as the implemented policy becomes more extreme. While a formal proof of this Lemma is presented in the Appendix A, the underlying reasoning is easy to illustrate graphically. In Figure 1, the solid lines represent the voter's expected period 1 payoff as a function of the implemented policy, 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 $-(x_1 - \underline{\alpha})^2$. The dashed curves instead represent the maximum and minimum possible values of the payoff realization when we take the random shock into account. Thus, the thick increasing dashed curves (representing the state of the world in which $x^v = \bar{\alpha}$) are, respectively, $-(x_1 - \bar{\alpha})^2 + \frac{1}{2\psi}$ and $-(x_1 - \bar{\alpha})^2 - \frac{1}{2\psi}$.

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 her true bliss point. 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. Clearly, 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. As the implemented policy becomes more extreme, the gray and black bullets get closer and closer to each other. The range of overlap becomes smaller, and the voter is more likely to learn the true value of her ideal policy.

Let me emphasize that the results presented below are robust to alternative assumptions about the distribution of the shock, as long as extreme policies are more informative than moderate ones. Consider for example a world in which the shock is normally distributed with full support. The learning process would be much smoother: any outcome realization 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 each signal's informativeness. Generally speaking, the mechanism that I uncover in this paper relies on the fact that the policy choice influences the amount of information the voter receives. This is what allows a dissenting ally to influence the equilibrium policy. As such, the main insights of the paper would survive in a world in which more moderate (rather than more extreme) policies are more informative.

The Incumbent

The voter's posterior beliefs determine her electoral decision, as shown in Lemma 1. Since the amount of information the voter observes is a function of the implemented policy, the incumbent has an incentive to engage in information control. The incumbent cannot control exactly which signal the voter will observe, but he can determine the expected probability of such a signal being informative. In other words, he cannot influence the voter's expected posterior (which is indeed always equal to the prior), but can influence its 'decomposition'. Hence, the first period equilibrium policy maximises the incumbent's trade-off between implementing his bliss point today and generating the optimal amount of information in order to get re-elected tomorrow. The way that this trade-off is optimised depends on the incumbent's ex-ante electoral strength. Define a *leading* incumbent as one who is guaranteed re-election if the voter receives no new information (condition (4) is satisfied at $\mu = \gamma$), and a *trailing* incumbent as one who will only be re-elected if the voter updates in his favour (condition (4) fails at $\mu = \gamma$ but is satisfied at $\mu = 1$). A *certain loser* is an incumbent who is replaced even if the voter updates in his favour (condition (4) fails to be satisfied at $\mu = 1$). The following Lemma holds:

Lemma 2: *In any PBE of the game*

- *A* ***certain loser*** *implements his bliss point*

  ($x_1^* = x^I$)

- *A* ***leading incumbent*** *implements a policy weakly more moderate than his bliss point*

  ($x_1^* \leq x^I$)

- *A* ***trailing incumbent*** *implements a policy weakly more extreme than his bliss point*

($x_1^* \geq x^I$)

For the incumbent, information revelation is risky. Even if $\gamma > \frac{1}{2}$, i.e. information is more likely to help him than hurt him, there is still a chance that the voter will instead learn that he own ideal policy is aligned with the challenger's (i.e. that $x^v = \underline{\alpha}$). A leading incumbent has no reason to accept the risk since he is guaranteed re-election when the voter does not update. Thus, he has incentives to prevent the voter from learning, and will always implement a policy that is (weakly) more moderate than his bliss point. Following Dewan and Hortala-Vallve (2017), I say that a leading incumbent experiences *fear of failure*. On the contrary, a trailing incumbent needs the voter to update (in his favour) in order to be re-elected. No matter how small the probability of success, a trailing incumbent always wants to engage in policy experimentation, so as to generate as much information as possible and improve its electoral prospects. Borrowing terminology from the IR literature (Downs and Rocke 1994), I say that this incumbent has incentives to *gamble for resurrection*, and always implements a policy (weakly) more extreme than his bliss point.[4] A certain loser trivially has no reason to engage in information control, since he cannot change the electoral outcome. Hence, he will always implement exactly his bliss point.[5] The exact policies adopted by a leading and a trailing incumbent are calculated in the Appendix. Such policies are a function of the bliss point x^I, the prior γ, and the probability of learning for any given policy ($4\alpha\psi$). The following Lemma defines the relationship between the equilibrium policy

[4] Notice that the incumbent's behaviour is reminiscent of the results in Groseclose (2001), despite the two models considering very different settings. In both papers a leading incumbent moderates in order to maximise his electoral advantage, while a trailing one moves to the extreme in order to exploit the variance in the distribution - of expected outcomes in this paper, of voters' bliss points in Groseclose.

[5] The same would apply to an incumbent is always re-elected, for all values of μ and $\mathbb{I}\delta$. However, given the symmetry assumption, such a case never occurs.

and the relevant parameters.

In Lemma 3 and the remainder of the paper I will be assuming that $x^I < x'$, where x' is the smallest (positive) policy that produces an informative signal with probability 1. The assumption is without loss of generality, and imposed in order to reduce the number of cases under consideration.

Lemma 3:

- *A trailing incumbent's equilibrium policy*
    1. *becomes (weakly) more extreme as his disadvantage decreases (γ increases)*
    2. *becomes (weakly) more extreme as his bliss point increases*
- *A leading incumbent's equilibrium policy*
    1. *becomes more extreme as his lead (γ) increases*
    2. *is always increasing in his bliss point when he enjoys a large lead ($\gamma > \frac{3}{4}$). When his lead is small ($\gamma < \frac{3}{4}$), the policy is non monotonic and concave in the bliss point*

The lower γ, the lower the probability that information will be in the incumbent's favor. As such, a leading incumbent's incentives to prevent information generation are stronger when γ is small, and a trailing incumbent's willingness to gamble is stronger when γ is large. Consider now the incumbent's bliss point. A trailing incumbent's policy choice is always increasing in his bliss point: as the incumbent becomes more extreme gambling becomes less costly and more valuable (since losing is more costly). Instead, a leading incumbent faces a trade off. As his bliss point becomes more extreme, preventing information generation becomes more costly today, but also more valuable for the future (as the challenger is further away and the payoff from winning increases). When the incumbent is too moderate, incentives to prevent information generation are weak since the gain from winning the next election is small. The direct effect dominates, and the equilibrium policy increases in the bliss point. Conversely, when the incumbent is too extreme, and the prior γ is sufficiently low, the electoral impact of the policy choice becomes dominant. As the incumbent's bliss point increases, winning the next election is more valuable, and the equilibrium policy becomes more moderate.

Dissent by an Extreme Ally

Moving one step back, we can now focus on the ally's decision whether to dissent against the incumbent. First of all, I establish that in equilibrium dissent is always harmful for the party's expected electoral performance, even if the incumbent best responds by modifying his policy choice precisely with the aim of minimizing this effect (as discussed in Lemma 3).

Lemma 4: *In equilibrium dissent always reduces the probability that the incumbent will be re-elected.*

Thus, by dissenting the ally reduces both his own and the incumbent's expected second period payoff. Nonetheless, dissent is sometimes observed in equilibrium. I show that, under some conditions, the ally faces a trade-off between maximizing the incumbent's electoral chances and inducing him to implement a policy more in line with his (i.e the ally's) own preferences.

The incumbent's ideological preferences are always more moderate than his ally's. However, by dissenting (and thereby harming the party in the upcoming election), the ally can induce him to implement a more extreme policy. Consider a leading incumbent. Absent dissent, he would always implement a policy (weakly) more moderate than his bliss point, in order to reduce the probability that the voter updates her beliefs about her true preferences. Suppose now that the incumbent's ally chooses to dissent against him. If the electoral cost is sufficiently large, this turns the leading incumbent into a trailing one. As Lemma 2 indicates, this creates incentives for the incumbent to gamble on resurrection: engage in extreme policies that increase the amount of voter learning. Thus, electorally costly dissent would move the incumbent's equilibrium policy choice to the extreme, closer to the ally's own preferences. When the gain is sufficiently large relative to the cost of losing the upcoming election, the ally chooses to dissent in equilibrium. Proposition 1 identifies necessary and sufficient conditions for this to occur (the proofs can be found in Appendix A). [6]

Proposition 1: *There exist $\underline{\gamma}$, $\overline{\gamma}$, $\underline{x^A}$ and $\underline{x^I}$ such that the incumbent's extreme ally chooses to dissent if and only if:*

- *Absent dissent, the incumbent is leading, but his advantage is not too large*
  $\underline{\gamma} < \gamma < \overline{\gamma}$, *where* $\underline{\gamma} \geq \frac{1}{2}$
- *The electoral cost of dissent is sufficiently high that it turns the leading incumbent into a trailing one, but not so high that the incumbent loses for sure*
  $(2\gamma - 1)4\bar{\alpha}x^I \leq \delta < 4\bar{\alpha}x^I$
- *Both the incumbent and his ally are sufficiently extreme*
  $x^I > \underline{x^I}$ and $x^A > \underline{x^A} > x^I$

The thresholds in Proposition 1 are a function of the other parameters in the model.

[6] In this paper I focus on dissent against an incumbent. However, the model can also be applied to explain the emergence of dissent within challenger parties. The challenger's ideological allies may want to openly attack him, thereby damaging the party in the upcoming election, so as to alter the incumbent's incentives to engage in information control. Within this framework, the challenger's allies use dissent to modify the incumbent's strategic choice, rather than to solve an ideological conflict within their own party. As such, dissent can emerge even absent any ideological disagreement, i.e. if the challenger and his allies have perfectly aligned preferences.

Intuition may suggest that dissent is more likely to materialize during periods of electoral crisis. The party is expected to perform poorly, and the ensuing internal turmoil degenerates into an open manifestation of conflict. The first result shows that, in the case of an extreme ally, the opposite is true. Suppose the incumbent is trailing even without experiencing dissent. Absent dissent, he will implement a policy that is weakly more extreme than his bliss point: he needs to generate information in order to be re-elected. Dissent either has no impact on his policy choice (if δ is so small that it does not affect the voter's electoral decision), or induces him to implement exactly his bliss point (if δ is sufficiently large to turn him into a sure loser). Hence, by dissenting the ally causes the incumbent to adopt a (weakly) more moderate policy, while also (weakly) reducing his own future expected payoff. Then, dissent is never observed in equilibrium. It is only when the incumbent is leading (i.e. $\gamma > \frac{1}{2}$) that the ally (potentially) gains from dissent by creating incentives to gamble for resurrection.

The second set of results refers to the electoral cost of dissent (δ). Quite intuitively, dissent never emerges in equilibrium when its electoral cost is so large that it makes the incumbent lose for sure. In this scenario the expected loss would be maximized, while the gain for the extreme ally would be minimized. Recall that an incumbent who is a sure loser has no reason to control information, and will always implement exactly his bliss point. Thus, while dissent would be somewhat effective in modifying the equilibrium policy, it could not induce the incumbent to move beyond his bliss point. The policy gain would be too small for the incumbent's ally to be willing to pay the cost of losing the upcoming election for sure. However, perhaps more surprisingly, the analysis also reveals that for dissent to be observed in equilibrium its electoral cost (δ) cannot be too small either. Recall that an incumbent is leading if the voter would choose to re-elect him upon receiving no new information. If δ is too small (relative to the prior γ), then the incumbent is still leading even after experiencing dissent. In this case, dissent has no effect on the voter's electoral choice and therefore no impact on the

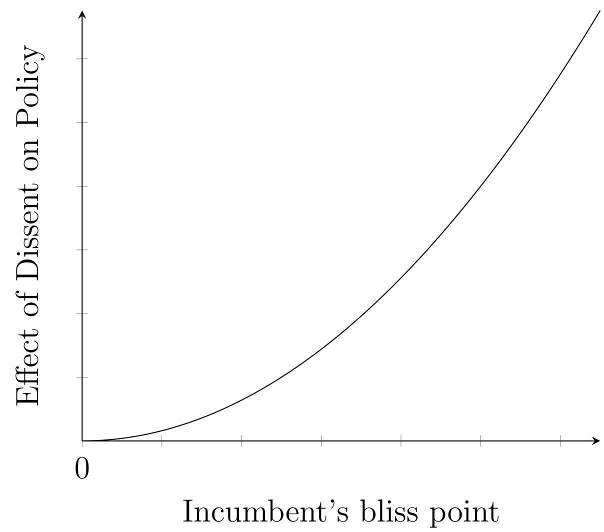

Figure 2: Effect of dissent on the equilibrium policy $((x_1^*|D=1)-(x_1^*|D=0))$

equilibrium policy. Trivially, the incumbent's ally has no reason to dissent in the first place. Thus, the electoral cost of dissent (δ) must be sufficiently large so as to turn a leading incumbent into a trailing one.

The second set of conditions on the prior γ ($\underline{\gamma} < \gamma < \overline{\gamma}$) ensures that the ally's gain from dissent outweighs the future expected cost. Recall that γ is the probability that the voter's true preferences are aligned with the incumbent's (i.e. $x^v = \overline{\alpha}$). As such, the higher γ, the lower a leading incumbent's incentives to prevent the voter from learning by implementing a moderate policy. As a consequence, the effect of dissent on the equilibrium policy is (weakly) decreasing in γ. When γ is too large (i.e. the incumbent enjoys a large lead) dissent therefore has a very small impact on the equilibrium policy, and the ally has no reason to pay the associated electoral cost. Conversely, if γ is too small the probability that the party would win the election after experiencing dissent is too low (recall that a trailing incumbent is re-elected only if the voter updates in his favor). Losing the upcoming election is very costly for the incumbent's extreme ally, therefore dissent is never observed in equilibrium.

Finally, let us now focus on the ideological misalignment between the incumbent and his allies. Such misalignment represents the only source of conflict in the model. Yet, Proposition 1 shows that increasing the ideological distance between the incumbent and his ally does not always make dissent more likely to emerge. Indeed, while only a sufficiently extreme ally may be willing to dissent, for this to occur in equilibrium the incumbent himself must be sufficiently extreme. Dissent cannot force the incumbent to implement any specific policy. The incumbent's ally can only influence his equilibrium choice by creating incentives to gamble on resurrection by engaging in policy experimentation. However, if the incumbent is too moderate, such incentives are too weak: gambling is too costly, and not very valuable. It is costly as it entails implementing extreme policies, potentially very far from the incumbent's bliss point. It is not very valuable since for a moderate incumbent the gain from winning the upcoming election is small (the distance from the opposition is small). Thus, as Figure 2 shows, the impact of dissent on the incumbent's choice is increasing in his bliss point (the vertical axis represents the difference between the equilibrium policy with and without dissent). If the incumbent is too moderate dissent will have a very small effect on the equilibrium policy. This reduces the ally's gain, and hence incentives to dissent in the first place. This result highlights the peculiar nature of dissent. in this model, which brings about unity even if no player actually cares about unity per se. Dissent serves the purpose of realigning the interests of the incumbent and his ally, thereby recomposing the existing ideological conflict. However, for dissent to be effective, such conflict cannot be too deep.

Comparative Statics: the Ambiguous Impact of the Ideological Conflict between the Incumbent and His Allies

Proposition 1 indicates that a necessary condition for dissent to occur in equilibrium is that γ falls within a certain range. The larger this range, the 'more likely' it is that dissent will be observed in equilibrium (in the sense of set inclusion). Proposition 2 describes how the size of this range (and therefore the likelihood of observing dissent) varies with the incumbent's and his ally's bliss points.

Proposition 2:

- *The likelihood of observing dissent (weakly) increases as the ally becomes more extreme*
- *There exists a unique* $\widehat{x^I}(x^A) > \underline{x^I}$ *such that if* $x^I < \widehat{x^I}(x^A)$*, then the likelihood of observing dissent increases as the incumbent becomes more extreme*

In line with the above discussion, Proposition 2 further highlights that the ideological conflict between the incumbent and his ally (i.e the distance between their bliss points) has an ambiguous effect on the probability of observing dissent. Increasing the ideological conflict either increases or decreases the likelihood of dissent, depending on whether the incumbent becomes more moderate or his ally more extreme. When the ideological misalignment increases due to the incumbent's ally becoming more extreme, dissent always becomes more likely. The more extreme the ally is, the more he gains by moving the equilibrium policy closer to his bliss point. However, the same is not necessarily true when the ideological conflict deepens due to the incumbent becoming more moderate. The intuition is exactly the same as discussed in relation to Proposition 1. As the incumbent becomes more extreme both a direct and indirect effects emerge. The direct effect is straightforward: the distance in the policy preferences of the incumbent and his ally decreases. This reduces the ally's incentives to dissent. The indirect effect goes in the opposite direction. As the incumbent becomes more extreme, dissent has a larger impact on his equilibrium policy choice. This in turn increases the ally's gain from dissent. If the incumbent's bliss point is sufficiently close to zero, this indirect effect dominates, and dissent is more likely to emerge as the ideological conflict decreases.

Welfare Analysis

Can the presence of the incumbent's extreme ally be welfare improving for the voter? In the model, the voter values policy experimentation as it increases the probability that she will make the correct electoral decision. As such, her first-period preferred platform maximizes the trade-off between her ex-ante ideological preferences (as dictated by her prior beliefs) and the need to learn about her ideal policy. However, the results presented above indicate that, under some conditions, electoral accountability may have the perverse consequence of inducing a lower lever of experimentation than what is optimal for the voter. The incumbent's extreme ally may mitigate such inefficiency, by inducing the incumbent to engage in extreme policies that increase the amount of voter learning. If the value of acquiring new information is sufficiently large, this strictly increases the voter's expected utility in the whole game. Proposition 3 identifies sufficient conditions for this to be true.

Proposition 3: *In equilibrium the voter benefits from the presence of an extreme ally to the incumbent party if:*

- *The cost of dissent δ is sufficiently large that it turns the leading incumbent into a trailing one, but not so large that it always hurts the voter ex ante ($\underline{\delta} < \delta < \overline{\delta_w}$)*
- *The value of information is sufficiently high*
  - *The prior (γ) is sufficiently close to $\frac{1}{2}$ ($\frac{1}{2} < \gamma < \overline{\gamma_w}$)*
  - *Incumbent and challenger are moderately polarized ($\underline{x^I_w} < x^I < \overline{x^I_w}$)*
  - *Learning the true state has a sufficiently large impact on the voter's preferences ($\bar{\alpha} > \bar{\alpha}_w$)*
- *The incumbent's ally is sufficiently extreme ($x^A > \underline{x^A_w}$)*

The first two conditions are intuitive: the voter must not dislike dissent too much, and obtaining new information must be sufficiently valuable. For this to be true, the voter's prior must be sufficiently uninformative (i.e. close to $\frac{1}{2}$), and the value of making the correct electoral decision must to be large enough. The third condition seems more puzzling: as the ally becomes more extreme the ideological misalignment with the voter increases. However, recall that the ally's bliss point has no direct effect on the equilibrium policy choice, thus on the voter's welfare. The effect is only an indirect one, through the ally's willingness to dissent. Since the first conditions impose further restrictions on the parameters, for the incumbent's ally to be willing to dissent when such conditions are satisfied (and therefore dissent is beneficial for the voter) he must be sufficiently extreme.

The normative implications of the results presented above are reminiscent of the 'case for responsible parties', presented by Bernhardt, Duggan and Squintani (2009). The authors find that, in a world in which the exact location of voters' preferences is unknown, all voters ex-ante prefer some degree of platform divergence between competing parties. However, electoral incentives may induce an excessive convergence in parties' platforms, thus ultimately hurting the voters. Therefore, as in this paper, a positive role of ideological extremism in the political elite emerges. In Bernhardt, Duggan and Squintani, all voters benefit from a moderate degree of parties' ideological extremism (polarization), that can guarantee an optimal level of platform divergence. In this paper, the incumbent's extreme ally can mitigate the perverse consequences created by electoral incentives, inducing an optimal level of policy experimentation.[7]

These results also speak to the debate on the normative evaluation of party factions.

[7] It is important to highlight that other mechanisms through which politicians' ideological polarization may prove welfare improving have also been identified. Van Weelden (2015) for example shows that polarization in the candidates' preferences decreases rent-seeking in equilibrium, unambiguously increasing voters' welfare for appropriate parameters.

The debate dates back to the 19th century. As noted by Boucek (2009), negative perceptions of factionalism originated with Hume (1877) and are still predominant. The main argument within this tradition is that factions 'exacerbate non-cooperative behaviour and so are antithetical to achievement of common goals' (Dewan and Squintani 2015, 861). A 'defence of factions' comes from the claim that organized and ideologically cohesive subgroups within political parties facilitate deliberation and pooling of valuable information, and therefore enhance the quality of the party's policy proposals. The argument is advanced initially by Bouceck herself (2009), investigated empirically by McAllister (1991), and proven formally by Dewan and Squintani (2015). The more or less implicit assumption is that factions engage in accommodative rather than disruptive activities (McAllister 1991). This paper moves one step further, showing that factionalism may have a positive value even when factions engage in 'disruptive' activities.

Micro Founding the Electoral Cost of Dissent

A key assumption of the paper is that dissent is electorally costly: everything else being equal, dissent reduces the probability that the incumbent is re-elected. In the model this cost is black-boxed, as this substantially simplifies both the analysis and interpretation of the results. However, it is worth discussing about potential ways to micro-found this assumption. Why do voters dislike parties that experience dissent?

One possibility is that dissent 'mechanically' reduces voters' appreciation of political parties. In this sense, the parameter δ would represent a behavioral 'bias' in the voters' preferences. It is well recognized that voters tend to like charismatic leaders (Groseclose 2001). Perhaps when the incumbent is publicly criticized by his own allies, or ridiculed by the media, this negatively affects voters' perception of the party.

However, dissent may be electorally costly even if voters are fully rational. Dissent may harm the party electorally because it conveys negative information to the voters. Voters do not dislike divided parties per se, rather the observation of dissent causes them to negatively update their beliefs over the incumbent's honesty, competence, etc. Dissent may convey such information in two different ways.

First, if the incumbent's allies have access to *verifiable* information, by dissenting they can expose him as a liar, corrupt or incompetent. The specification and results of this model would be exactly as presented above, with δ representing the electoral value of competence (net of the probability that the challenger is a good type). Under the conditions identified in Proposition 1, the allies choose to dissent whenever they can reveal evidence that the incumbent is a bad type. If the conditions are not met, the allies always keep quiet. The only difference with the model presented here is that, because verifiable information cannot be fabricated, dissent can never emerge in equilibrium if the incumbent is a good type.

Alternatively, we may assume that the incumbent's allies do not have access to such verifiable evidence. Nonetheless, they may have an informational advantage with respect to the voters. For example, the allies may scrutinize the incumbent's previous actions and performance, thereby obtaining additional information about his true competence (see Caillaud and Tirole 1999, Fox and Van Weelden 2010). As such, the allies can engage in a signaling game with the electorate. Dissent is electorally costly when, *in equilibrium*, it constitutes a negative signal of the incumbent's type. However, for dissent to emerge when it is electorally costly (i.e. under separation or semi-separation), the gain from changing the incumbent's policy choice must be sufficiently large. The qualitative results would then be as in the reduced-form model presented above.[8]

In concluding this section let me emphasize that, while the assumption of electorally

[8]Some additional conditions are however required to sustain separation. The ally must care sufficiently about quality, and he cannot be too extreme, as otherwise he would have incentives to dissent even when the incumbent is a good type.

costly dissent is motivated by both empirical evidence and the theoretical reasoning presented above, the mechanism identified in this paper relies only on the voters not being indifferent to dissent (i.e. $\delta \neq 0$). Indeed, it would survive in a world in which dissent produces a positive valence shock (i.e. $\delta < 0$), thus improving rather than damaging a party's electoral prospects. Clearly, under such an assumption the puzzle would be reversed: if dissent is beneficial to a party, how do we explain cases in which we do not observe dissent? The mechanism identified in this paper would provide a potential answer. An extreme ally may choose not to improve its party's electoral prospects, in order to preserve the incumbent's incentives to gamble for resurrection.

Extension: What if the Ally Has Bargaining Power?

So far I have assumed that the incumbent is essentially a policy dictator. His allies have no formal bargaining power, and dissent is the only tool to influence the equilibrium policy. This assumption is plausible if we consider dissent by ideological allies external to the party (such as ideologically aligned media), but perhaps less so if we focus on factional dissent. To be sure, institutional features such as the vote of confidence procedure (in parliamentary systems) may grant large discretion to the party leader. However, it is also possible that the minority faction will have some bargaining power (perhaps due to the credible threat of a formal defection/party split) and therefore influence over policy.

Are the results of this paper robust to assuming that the incumbent's allies have bargaining power over the policy making process? To address this question, I analyse the model under the assumption that, in the first period, the incumbent maximises a weighted average of his own and his ally's utility:

$$U_1^W = (1-\beta)U^I + \beta U^A \tag{5}$$

This is equivalent to considering (in a reduced form) a game in which *after* the incumbent's ally chooses whether to dissent against him, the two engage in in a bargaining game over the policy choice. β thus represents the ally's influence over policy making in the first period. There are two reasons to consider a setting in which the ally has bargaining power in the first period only. First, it is plausible to argue that faction's bargaining power comes from the threat of a formal defection. Such a threat is credible when the ally (i.e. the dissenting faction) has a sufficiently high chance of winning the upcoming election if running alone after a party split (as for example in Mutlu-Eren 2015; Hortala-Vallve and Mueller 2015). Since there is no election after the second period, the ally has no way to make a credible threat. Additionally, assuming that the ally has bargaining power only over the first period policy is a way to obtain a meaningful comparison with the baseline model. Suppose that the second period policy is also determined via a bargaining process. Recall that dissent occurs in equilibrium only if the incumbent is leading. In the baseline model this requires $\gamma > \frac{1}{2}$ (since incumbent and challenger are assumed to be symmetric). If the extreme ally has formal bargaining power over the second period policy, the condition becomes $\gamma > \frac{(\beta x^A + (1-\beta)x^I + \alpha)^2 - (x^I - \alpha)^2}{4\alpha(x^I + \beta x^A + (1-\beta)x^I)} > \frac{1}{2}$. Therefore, when comparing the bargaining extension to the baseline model, I would not only be altering the β parameter, but also imposing further conditions on γ, which would make the comparison less meaningful.

It is straightforward to see why bargaining power and dissent are, to a certain extent, substitutes. Dissent is a tool to influence the incumbent's equilibrium policy choice. When the ally has some formal control over policy making ($\beta > 0$), the incentives to pay the electoral cost of dissent are weaker. Indeed, in the limiting case in which the ally is given full authority over policy ($\beta = 1$), he will never choose to dissent against himself.

However, the results uncover a second and more subtle effect. Bargaining power and dissent will sometimes complement rather than substitute each other, so that dissent is more likely to be observed compared to the case in which the incumbent is a policy dictator. Recall that, in the no-bargaining baseline, dissent emerges in equilibrium only if the incumbent is sufficiently extreme; if the incumbent is too moderate, the incentives to gamble are too weak and dissent has too little an effect on the equilibrium policy. However, if the ally is given formal authority over policy making, it can effectively 'compensate' for an excessively moderate incumbent, so that dissent can emerge in equilibrium for every (positive) value of the incumbent's bliss point. The following holds:

Proposition 4: *For all $x^I \geq 0$, there exist non-measure zero sets $\Gamma(x^I)$ and $B(x^I)$ such that if $\gamma \in \Gamma(x^I)$ and $\beta \in B(x^I)$ then dissent by an extreme ally occurs in equilibrium*

The sets $B(x^I)$ and $\Gamma(x^I)$ also depend on the other parameters. Proposition 4 shows that the mechanism uncovered in the paper is robust to assuming that the ally has formal bargaining power, arising for example from a credible threat of defection or party split. Additionally – as discussed above – when the incumbent is sufficiently moderate (x^I is sufficiently close to zero) $\beta > 0$ is a necessary condition to observe dissent in equilibrium. The following corollary also holds:

Corollary 1: *Suppose that* $\frac{1}{8\overline{\alpha}\psi} < x^I$ and $\frac{1}{4\overline{\alpha}\psi} < x^A < \frac{1}{4\overline{\alpha}\psi(1-2\overline{\alpha}\psi x^I)}$. *Then, for all $\beta \in [0,1)$, there exists a non-measure zero set $\Gamma(\beta)$ such that if $\gamma \in \Gamma(\beta)$ dissent by an extreme ally occurs in equilibrium*

The corollary shows that even if the ally is granted almost full discretion over the first period policy (β approaches 1), he will still choose to dissent under some conditions.

Empirical Implications and Falsifiability

Before concluding, it is important to explore the theory's implications for empirical research, and to discuss how the model's predictions may allow us to adjudicate between competing explanations for the emergence of dissent.

Several scholars have recognized that, as highlighted in this paper, party unity may have a crucial impact on electoral outcomes. Trying to quantify the electoral cost of dissent is therefore an important empirical exercise, useful to complete our understanding of electoral competition and accountability. The strategy employed in the extant literature is to regress the probability of winning (or other measures of electoral success) at time t on a binary variable indicating whether the party experienced dissent at time $t-1$ (e.g. Clark 2009, Kam 2009, Groeling 2010):

$$prob(W_i = 1) = \alpha + \beta_1 X_i + \beta_2 D_i + \epsilon_i \tag{6}$$

Where X_i is a vector of covariates, and β_2 is the coefficient of interest. Graphically, the quantity of interest is the average distance between the two curves in Figure 3, representing the probability of winning as a function of the party's ex-ante electoral strength (γ), with and without dissent.

The results of the model have two key implications. First of all, they show that it is impossible to isolate the direct effect, i.e. voter's dislike of parties that experience dissent. The incumbent modifies his policy choice precisely with the aim of mitigating the electoral cost of dissent. Thus, any estimate would at best reflect the equilibrium effect of dissent on electoral success: the cost mediated through the incumbent's best response. Additionally, the model shows that any such estimate would inevitably suffer from selection bias. Proposition 1 shows that whether parties experience dissent depends precisely on their ex-ante electoral strength (γ needs to be moderately high). Thus, it is impossible to observe both treated and control units for the same level of

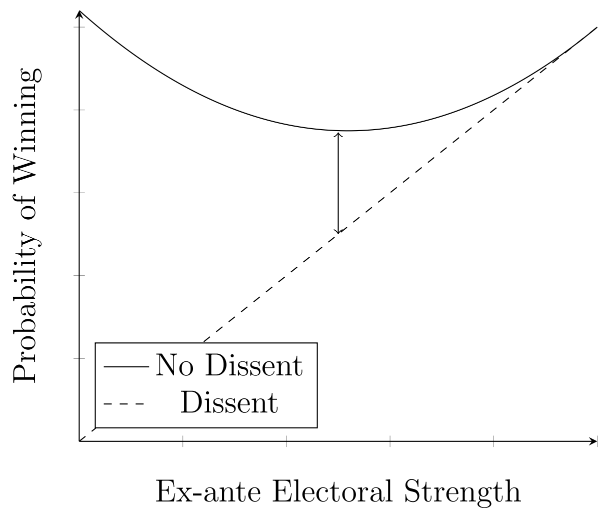

Figure 3: Probability of Winning as Function of Ex-ante Strength (γ).
The blue line represents the probability of winning when the incumbent experiences no dissent. The red line represents the probability of winning conditional on experiencing dissent.

ex-ante electoral strength. Figure 4 represents what the researcher can actually observe: treated units, i.e. parties that experience dissent, at moderately high levels of electoral strength and untreated ones at γ close to $\frac{1}{2}$ and 1. Comparing parties that experience dissent with their untreated counterparts means comparing parties with different levels of ex-ante electoral strength. Thus, it is impossible to recover an unbiased estimate of the (equilibrium) effect of dissent on parties' electoral performance.

Further, it is hard to know ex-ante what the direction of the bias in the results will be. In the example of Figure 4, the direction of the bias would be upward: the estimated electoral cost of dissent would be higher than the true one. However, under different parameter values, the dissenting region shifts. Consider for example Figure 5, obtained by increasing the ally's bliss point: dissent emerges only at higher values of γ (compared to the case illustrated in Figure 4). In this case the direction of the bias is no longer clear. Indeed, the estimate may even have the wrong sign. Due to the selection bias, parties that experience dissent may perform better than 'control' units. Thus, even if we are aware of the existence of the bias, it is hard to interpret the results

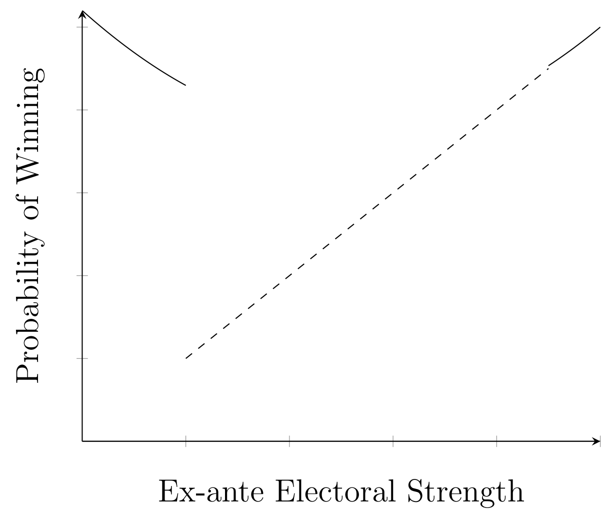

Figure 4: Probability of Winning as Function of Ex-ante Strength (γ) - Observable

of this type of analysis.

However, this does not imply that the theory is not falsifiable. Indeed, the model generates predictions regarding the electoral performance of parties that do experience dissent. If we focus on this restricted sample, thereby avoiding the problem of selection bias described above, we can still say something on when dissent is expected to harm parties the most.

When a party experiences dissent, winning the upcoming election requires the voter to discover her true preferences (and update in favour of the party). The probability of the voter learning is obviously higher in high information environments. The larger the amount of information received by the voters, and their ability to interpret such information, the higher the probability of winning conditional on experiencing dissent. Thus, when considering a regression with the probability of winning conditional on dissent as dependent variable, we should expect to see a positive and statistically significant coefficient for variables such as news media consumption, education or political engagement in the population. Additionally, irrespective of the noisiness of the information environment, the incumbent must be willing to gamble and engage in policy experimentation. Thus, conditional on experiencing dissent, we should expect the party's electoral suc-

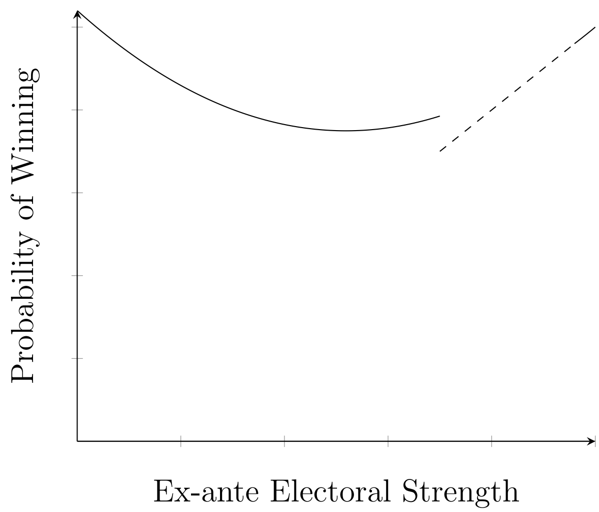

Figure 5: Probability of Winning as Function of Ex-ante Strength, $x^A > \underline{x^A}$

cess to be increasing in the leader's ideological extremism. The more extreme the party leader is, the more he will be willing to gamble (gambling is both less costly and more valuable), the more likely it is that the policy outcome will be informative.

Finally, it is important to discuss whether the theory's empirical implications may allow us to distinguish it from alternative explanations for the emergence of electorally costly dissent.One possibility is that dissent emerges when politicians face a trade-off between the national party's electoral fortune and their individual success. This theory, and the differences with the argument presented here, have already been discussed in the literature review. Here, it is important to add that the mechanism identified in this paper applies not only to intra-party dissent but also to cases in which the party leader is attacked by ideological allies external to the party (such as media outlets, as discusses in the introduction). The same does not seem to hold for the explanation relying on the trade off between individual and collective reputation.

Another possibility is that the dissenters are trying to damage the leader so as to make it easier to depose him. Within this framework, dissent should emerge when the leader is expected to perform poorly in the upcoming elections. Yet, this is not always the case. In the Italian example, dissent exploded against a leader who was expected to bring the party to electoral success. This is in line with the predictions of Proposition 1, according to which dissent by an extreme ally emerges when the incumbent is leading and expected to win with a sufficiently high probability. Finally, it is important to stress that if the dissenters' goal is to replace the dominant faction and take over the party (rather than simply depose the incumbent leader), then this argument complements the one proposed in this paper. For an extreme faction to take over, it has to believe that it has a chance of winning the election. When the electorate is too moderate, this requires changing voters' policy preferences. That is, the faction has incentives to force the incumbent leader to experiment, just like in the present model. As such, the framework presented here could explain dissent for pure policy-motivated reasons (as in the current paper) or for both policy *and* instrumental reasons (taking over). In this perspective allowing for replacement does not alter (and if anything strengthens) the qualitative insights presented here.

Conclusion

Political leaders often experience dissent by their own ideological allies despite this being electorally harmful. In order to address this puzzle, I have presented a political agency model in which voters are learning about their own policy preferences over time. The first contribution of the paper is to provide a new framework to study policy experimentation. Within this framework, the amount of voter learning depends on the location of the implemented policy, with extreme policies generating more information. As a consequence, leading incumbents have incentives to implement moderate platforms, while trailing ones want to engage in extreme policies that generate more information.

Within this setting, dissent emerges precisely because it is electorally costly. By dissenting (and thereby harming the party in the upcoming election) the incumbent's ally can change his incentives to choose more or less extreme policies, which affects the amount of voter learning. This creates a trade-off between winning the upcoming election and inducing the incumbent to pursue the ally's all-things-considered more-preferred policy. Optimally balancing this trade-off sometimes involves active dissent that damages the party in the short-run. We observe dissent in equilibrium when the ensuing electoral cost is sufficiently high, the incumbent's bliss point is sufficiently biased in the direction of his ally, and the voter's prior that her preferences are in line with the incumbent's is moderately high. The results are robust to assuming that the allies have formal bargaining power over the first period policy. Further, the results indicate that the presence of extremists within the incumbent party may be welfare improving for the voter, inducing an optimal amount of policy experimentation by the office holder. The theory also has relevant implications for empirical research, showing that existing estimates of the electoral rewards of party unity obtained by comparing treated and control units are inevitably biased. However, the model generates testable predictions regarding parties' electoral performance conditional on dissent.

In this paper I have assumed that an ideological conflict underlies the emergence of this form of dissent. However, within the same framework dissent can arise even absent any ideological disagreement. For example, the incumbent and his ally may have the same bliss point but different access to (or evaluation of) office rents. This induces different risk appetite in policy making, thereby potentially generating a conflict in preferences. Similarly, the two actors may disagree about their beliefs over the voters' ideal policy. Suppose, for example, that the ally assigns a higher probability to the voter's true preferences being aligned with the party's. Then, the ally would always prefer a (weakly) more extreme policy relative to the incumbent. As in the model presented here, electorally costly dissent would therefore serve the purpose of incentivizing the incumbent to engage in policy experimentation, (re)creating unity of interests with his ideological allies.

Finally, while this paper has focused on intra-party conflict, the mechanism it uncovers applies more generally. Indeed, it can capture the dynamics of the interaction between political actors in any strategic situation that can be described as a principal agent model with two key features. First, there is some (common) uncertainty on what is the principal's optimal retention decision, and the amount of information that is generated is a function of the agent's action. The principal's uncertainty can refer to her ideal policy, as in the model presented here, or to the agent's type, e.g. his ability or competence. Second, the agent's ally (i.e. an actor whose payoff is higher when the agent is retained than when he is replaced) can take an action that, everything else constant, changes the probability that the agent is retained. In this setting, the ally can choose to manipulate the probability that the agent is retained in order to alter his incentives to generate more or less information. This creates a trade off between ensuring that the agent is retained, and inducing him to take an action closer to the ally's own preferences. The agent's action may refer to a level of effort, a point in the policy space, the degree of reform, or even the amount of rent extraction. As such, this framework can be applied to several different settings, encompassing both developed and developing democracies, as well as authoritarian countries.

References

[1] Austen-Smith, David. 1998. "Allocating Access for Information and Contributions" *Journal of Law, Economics, and Organization* 14:2 :277-303.

[2] Belloni, Frank P., and Dennis C. Beller. 1976. "The study of party factions as competitive political organizations" *Western Political Quarterly*, 29:4: 531-549.

[3] Bernhardt, Dan, John Duggan, and Francesco Squintani. 2009. "The case for responsible parties" *American Political Science Review*, 103:4: 570-587.

[4] Bersani, Pierluigi. 2015. "Renzi? Governa col mio 25%" *Il Giornale*

[5] Boucek, Francoise. 2009. "Rethinking factionalism: typologies, intra-party dynamics and three faces of factionalis" *Party Politics* 15:4: 455-485.

[6] Budge, Ian, Lawrence Ezrow, and Michael D. McDonald. 2010. "Ideology, Party Factionalism and Policy Change: an Integrated Dynamic Theory" *British Journal of Political Science* 40:4: 781-804.

[7] Buisseret, Peter and Carlo Prato. 2018. 'Competing Principals? Legislative Representation in List PR Systems'. Working Paper.

[8] Caillaud, Bernard, and Jean Tirole. 1999. "Party governance and ideological bias" *European Economic Review* 43:4: 779-789.

[9] Callander, Steven. 2011. "Searching for Good Policies" *American Political Science Review* 105:4: 643-662.

[10] Callander, Steven and Patrick Hummel. 2014. "Preemptive policy experimentation" *Econometrica* 82: 4: 1509-1528.

[11] Clark, Michael. 2009. "Valence and electoral outcomes in Western Europe, 1976-1998". *Electoral Studies* 28: 1: 111-122.

[12] Coleman, J. S. (1971), 'Internal processes governing party positions in elections', *Public Choice*, 11: 1: 35-60.

[13] Cuperlo, Giovanni. 2016. "Cuperlo contro Renzi: non sei all'altezza e sei un capo arrogante" *Il Messaggero*

[14] D'Alema, Massimo. 2016. "Pur di cacciare Renzi, sono pronto a votare anche Raggi" *La Repubblica*.

[15] Dewan, Torun and Francesco Squintani. 2015. "In defense of factions" *American Journal of Political Science* 60:4: 860-881.

[16] Dewan, Torun and Rafael Hortala-Vallve. 2017. "Electoral competition, control and learning" *British Journal of Political Science* 1-17.

[17] Downs, George W., and David M. Rocke. 1994. "Conflict, agency, and gambling for resurrection: The principal-agent problem goes to war" *American Journal of Political Science* 38:2: 362-380.

[18] Edwardes, Charlotte. 2017. "Battle of the Tories: who will succeed Theresa May?" *The Evening Standard*.

[19] Fox, Justin and Richard Van Weelden. 2010. "Partisanship and the effectiveness of oversight" *Jourmal of Public Economics* 94: 674-687.

[20] Greene, Zachary David, and Matthias Haber. 2015. "The consequences of appearing divided: An analysis of party evaluations and vote choice" *Electoral Studies* 37: 15-27.

[21] Groeling, Tim. 2010. *When politicians attack: Party cohesion in the media* Cambridge University Press.

[22] Groseclose, Tim. 2001. "A model of candidate location when one candidate has a valence advantage" *American Journal of Political Science*, 45: 4: 826-866.

[23] Hortala-Vallve, Rafael and Hannes Mueller. 2015. "Primaries: the unifying force" *Public Choice* 163:3: 289-305.

[24] Jones, Owen. 2016. "Jeremy Corbyn's critics must decide: unity or terminal decline" *The Guardian.*

[25] Kam, Christopher J. 2009. *Party discipline and parliamentary politics* Cambridge University Press.

[26] Kamenica, Emir, and Matthew Gentzkow. 2011. "Bayesian persuasion" *The American Economic Review* 101: 6: 2590-2615.

[27] Majumdar, Sumon, and Sharun W. Mukand. 2004. "Policy gambles" *The American Economic Review* 94: 4: 1207-1222.

[28] Mandelson, Peter. 2017. "I try to undermine Jeremy Corbyn 'every single day'" *The Guardian.*

[29] McAllister, Ian 1991. "Party Adaptation and Factionalism within the Australian Party System" *American Journal of Political Science*, 35:3: 206-227.

[30] McCain, John. 2017. Speech filmed at Munich Security Conference in Germany. $https://www.washingtonpost.com/news/the-fix/wp/2017/02/17/john-mccain-just-systematically-dismantled-donald-trumps-entire-worldview/?utm_term=.15a944873049$

[31] Mutlu-Eren, Hande. 2015. "Keeping the party together" *Public Choice* 164:1: 117-133.

[32] Sartori, Giovanni. 1976. *Parties and party systems: A framework for analysis* CUP Archive.

[33] Tonybee, Polly. 2016. "Dismal, lifeless, spineless - Jeremy Corbyn let us down again" *The Guardian.*

[34] Urwin, Rosamund. 2017. "The next generation - do government leaders have a 'political shelf life'?" *The Evening Standard*.

[35] Van Weelden, Richard. 2015. "The welfare implications of electoral polarization" *Social Choice and Welfare* 45: 653-686.

Appendix A: Proofs

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

(i) Her posterior μ takes one of three values: $\mu \in \{0, \gamma, 1\}$;

(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 $\mu \neq \gamma$ 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.

Symmetric results apply when $x_t < 0$.

Proof. The proof of part (i) is trivial given the boundedness of the distribution of ϵ, 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 (ϵ): $U_t^v = -(x^v - x_t)^2 + \epsilon$. Denote as $f(\cdot)$ the PDF of ϵ. Then,

$$prob(x^v = \bar{\alpha}|U_t^v) = \frac{f(U_t^v + (x_t - \bar{\alpha})^2)\gamma}{f(U_t^v + (x_t - \bar{\alpha})^2)\gamma + f(U_t^v + (x_t - \underline{\alpha})^2)(1 - \gamma)} \quad (7)$$

Given the assumption that ϵ is uniformly distributed

$$f(U_t^v + (x_t - \bar{\alpha})^2) = f(U_t^v + (x_t - \underline{\alpha})^2) \tag{8}$$

Therefore the above simplifies to

$$prob(x^v = \bar{\alpha}|U_t^v) = \gamma \tag{9}$$

This concludes the proof of Claim 1. □

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 $|x_1| \geq |x'|$*

$$Prob(L = 1|x_1) = 1 \tag{10}$$

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

$$Prob(L = 1|x' \geq x_1 \geq 0) = 4\bar{\alpha}\psi x_1 \tag{11}$$

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

$$Prob(L = 1| - x' \leq x_1 \leq 0) = -4\bar{\alpha}\psi x_1 \tag{12}$$

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

$$-(x_t-\underline{\alpha})^2 + \frac{1}{2\psi} + (x_t-\bar{\alpha})^2 + \frac{1}{2\psi} \leq 0 \tag{13}$$

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

$$x \geq \frac{1}{4\bar{\alpha}\psi} = x' \tag{14}$$

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

$$Prob(L=1|\gamma, 0 < x_1 < x') =$$

$$\gamma[Prob(-(x_t-\bar{\alpha})^2+\epsilon_1 > -(x_t-\underline{\alpha})^2+\tfrac{1}{2\psi})] + (1-\gamma)[Prob(-(x_t-\underline{\alpha})^2+\epsilon_1 < -(x_t-\bar{\alpha})^2-\tfrac{1}{2\psi})] \tag{15}$$

Given the symmetry

$$Prob(-(x_t-\bar{\alpha})^2+\epsilon_1 > -(x_t-\underline{\alpha})^2+\frac{1}{2\psi}) = Prob(-(x_t-\underline{\alpha})^2+\epsilon_1 < -(x_t-\bar{\alpha})^2-\frac{1}{2\psi}) \tag{16}$$

(15) simplifies to

$$Prob(L=1|x_1>0) = Prob(-(x_t-\bar{\alpha})^2+\epsilon_1 > -(x_t-\underline{\alpha})^2+\frac{1}{2\psi})) = 4\bar{\alpha}\psi x_1 \tag{17}$$

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

This concludes the proof of Claim 2 □

This concludes the proof of Lemma 1 □

In what follows I will assume that $x^I < \frac{1}{4\bar{\alpha}\psi}$. This assumption is without loss of generality, and imposed in order to reduce the number of cases under consideration; results for the case in which $x^I > \frac{1}{4\bar{\alpha}\psi}$ are available upon request.

Lemma 2: *In any PBE of the game*

- *A* ***certain loser*** *implements his bliss point*
  ($x_1^* = x^I$)
- *A* ***leading incumbent*** *implements a policy weakly more moderate than his bliss point*
  ($x_1^* \leq x^I$)
- *A* ***trailing incumbent*** *implements a policy weakly more extreme than his bliss point*
  ($x_1^* \geq x^I$)

Proof. The proof of the first point is trivial: a certain loser is never re-elected, hence his policy choice does not influence his future payoff. He maximises his immediate utility by implementing his bliss point x^I. Leading and trailing incumbents will instead consider the expected informativeness of the policy, and how it influences their probability of re-election.

Consider first a trailing incumbent. Electoral concerns create incentives to implement an informative policy. The equilibrium policy solves the following maximisation problem:

$$\begin{aligned} \underset{x_1}{\text{maximise}} \quad & -(\mathrm{x_1} - \mathrm{x^I})^2 - (1 - 4\bar{\alpha}\psi \mathrm{x_1}\gamma)(\mathrm{x^I} + \mathrm{x^I})^2 - 4\bar{\alpha}\psi \mathrm{x_1}\gamma(\mathrm{x^I} - \mathrm{x^I})^2 \\ \text{subject to} \quad & x_1 \leq \frac{1}{4\bar{\alpha}\psi} \end{aligned} \tag{18}$$

Hence:

$$x_1^* = \min\{x^I + 8\bar{\alpha}\psi(x^I)^2\gamma, \frac{1}{4\bar{\alpha}\psi}\} \tag{19}$$

The condition that $x_1 \leq \frac{1}{4\alpha\psi}$ derives from the fact that any policy weakly more extreme than $x' = \frac{1}{4\alpha\psi}$ is fully informative, therefore the leading incumbent would have no reason to move beyond x' (recall that we assume $x^I < x'$).

Consider now a leading incumbent. The equilibrium policy will maximise the trade-off between implementing his true bliss point today, and generating as little information as possible, so as to increase the probability of being re-elected tomorrow. The equilibrium policy solves the following maximisation problem:

$$\begin{aligned} \underset{x_1}{\text{maximise}} \quad & -(\mathrm{x_1} - \mathrm{x^I})^2 - 4\bar{\alpha}\psi \mathrm{x_1}(1-\gamma)(\mathrm{x^I} + \mathrm{x^I})^2 - (1 - 4\bar{\alpha}\psi \mathrm{x_1}(1-\gamma))(\mathrm{x^I} - \mathrm{x^I})^2 \\ \text{subject to} \quad & x_1 \geq 0 \end{aligned} \tag{20}$$

Hence:

$$x_1^* = \max\{x^I - 8\bar{\alpha}\psi(x^I)^2(1-\gamma), 0\} \tag{21}$$

An incumbent may only be leading if $\gamma > \frac{1}{2}$. Additionally, $x^I < \frac{1}{4\bar{\alpha}\psi}$ by assumption. Thus, $x^I - 8\bar{\alpha}\psi(x^I)^2(1-\gamma) > 0$ and

$$x_1^* = x^I - 8\bar{\alpha}\psi(x^I)^2(1-\gamma) \tag{22}$$

□

Lemma 3:

- *A trailing incumbent's equilibrium policy*
  1. *becomes (weakly) more extreme as his disadvantage decreases (γ increases)*
  2. *becomes(weakly) more extreme as his bliss point increases*
- *A leading incumbent's equilibrium policy*
  1. *becomes more extreme as his lead (γ) increases*
  2. *always becomes more extreme as his bliss point increases, when he enjoys a large lead ($\gamma > \frac{3}{4}$). When the lead is moderate ($\gamma < \frac{3}{4}$), the policy is non monotonic and concave in the bliss point*

The proof is omitted since it follows straightforwardly from Lemma 2.

Lemma 4: *In equilibrium dissent always (weakly) reduces the probability that the incumbent will be re-elected.*

Proof. Let x^d be the incumbent's policy choice after dissent, and x the policy that he would choose otherwise. Consider first of all a leading incumbent. We must distinguish between three cases: (i) $\delta < \underline{\delta}$ such that the incumbent's initial advantage is not outweighed (i.e. $\gamma > \frac{\delta + 4\bar{\alpha}\psi x^I}{8\bar{\alpha}\psi x^I}$). In this case dissent does not modify the incumbent's policy choice nor the voter's electoral decision. (ii) $\delta \geq \bar{\delta}$ such that the incumbent always loses after dissent (i.e $\delta \geq 4\bar{\alpha}x^I$). The claim follows straightforwardly. (iii) $\delta \in [\underline{\delta}; \bar{\delta})$, such that the incumbent wins if and only if the voter updates in his favour (i.e. dissent turns the leading incumbent into a trailing one). The following holds. Let $\pi(x_1)$ be the probability of the voter observing an informative signal at the end of period 1, as a function of the implemented policy. The probability of the incumbent being re-elected absent dissent is $1-\pi(x)(1-\gamma)$. The probability of the incumbent being re-elected after dissent is instead $\pi(x^d)\gamma$. $1-\pi(x)(1-\gamma) \geq \pi(x^d)\gamma$, since the LHS is at least $1-(1-\gamma)=\gamma$ and the RHS is at most γ. Finally, consider a trailing incumbent. There are only two possibilities: (i) $\delta > \underline{\delta}$ such that after experiencing dissent the incumbent loses for sure. The claim follows trivially (ii) $\delta \leq \underline{\delta}$ such that the incumbent is still trailing even after experiencing dissent. Dissent has no impact on the policy choice, nor on the voter's electoral decision. □

Proposition 1: *There exist $\underline{\gamma}$, $\overline{\gamma}$, $\underline{x^A}$ and $\underline{x^I}$ such that the incumbent's extreme ally chooses to dissent if and only if:*

- *Absent dissent, the incumbent is leading, but his advantage is not too large*
  $\underline{\gamma} < \gamma < \overline{\gamma}$, *where* $\underline{\gamma} \geq \frac{1}{2}$
- *The electoral cost of dissent is sufficiently high that it turns the leading incumbent into a trailing one, but not so high that the incumbent loses for sure*
  $(2\gamma-1)4\bar{\alpha}x^I \leq \delta < 4\bar{\alpha}x^I$
- *Both the incumbent and his ally are sufficiently extreme*
  $x^I > \underline{x^I}$ and $x^A > \underline{x^A} > x^I$

Proof. Let me first prove that dissent is never observed in equilibrium if the incumbent is trailing.[9] Absent dissent, a trailing incumbent implements policy $x_1^* = \min \in \{x^I + 8\bar{\alpha}\psi(x^I)^2\gamma, \frac{1}{4\bar{\alpha}\psi}\}$. Dissent has no impact on his policy choice (and therefore never

[9]The reason why we must consider this case is that, even if dissent would always induce a trailing incumbent to moderate his policy choice, it may be the case that the ally's bliss point lies between the platforms that the incumbent would implement with and without dissent. It is therefore possible that $[x^A - (x_1^*|D=1, \gamma < \frac{1}{2})]^2 < [x^A - (x_1^*|D=0, \gamma < \frac{1}{2})]^2$ even if $(x_1^*|D=1, \gamma < \frac{1}{2}) < (x_1^*|D=0, \gamma < \frac{1}{2})$. We must therefore exclude that the ally's gain from dissenting against a trailing incumbent is larger than the cost.

emerges in equilibrium) if $\delta < 4\bar{\alpha}x^I$. Suppose instead that $\delta \geq 4\bar{\alpha}x^I$. Then, after experiencing dissent the incumbent is a sure loser: even if the voter learns that $x^v = \bar{\alpha}$, she will still choose to replace the incumbent with his challenger. As a consequence, in the first period the incumbent would always implement exactly his bliss point x^I upon experiencing dissent.

Therefore, there are two pairs of equilibrium policies that we must consider:

- $(x_1^*|D=1) = x^I$ and $(x_1^*|D=0) = \frac{1}{4\bar{\alpha}\psi}$, when $x^I + 8\bar{\alpha}\psi(x^I)^2\gamma \geq \frac{1}{4\bar{\alpha}\psi}$
- $(x_1^*|D=1) = x^I$ and $(x_1^*|D=0) = x^I + 8\bar{\alpha}\psi(x^I)^2\gamma$, when $x^I + 8\bar{\alpha}\psi(x^I)^2\gamma < \frac{1}{4\bar{\alpha}\psi}$

I will analyse each case separately, conjecturing the existence of an equilibrium in which the ally chooses to dissent.

Case 1: $(x_1^*|D=1) = x^I$ and $(x_1^*|D=0) = \frac{1}{4\bar{\alpha}\psi}$

The equilibrium conditions for the incumbent are

$$\gamma \leq \frac{1}{2} \tag{23}$$

$$\gamma > \frac{1}{8\bar{\alpha}\psi x^I}\left(\frac{1}{4\bar{\alpha}\psi x^I} - 1\right) \tag{24}$$

Additionally, the equilibrium condition for the faction is

$$-(x^I - x^A)^2 - (x^I + x^A)^2 > -\left(\frac{1}{4\bar{\alpha}\psi} - x^A\right)^2 - \gamma(x^I - x^A)^2 - (1-\gamma)(x^I + x^A)^2 \tag{25}$$

Which reduces to

$$\gamma < \frac{-8\bar{\alpha}\psi x^A(1 - 4\bar{\alpha}\psi x^I) - (4\bar{\alpha}\psi x^I)^2 + 1}{(8\bar{\alpha}\psi)^2 x^I x^A} \tag{26}$$

Thus, we need to identify conditions under which

$$\frac{-8\bar{\alpha}\psi x^A(1-4\bar{\alpha}\psi x^I)-(4\bar{\alpha}\psi x^I)^2+1}{(8\bar{\alpha}\psi)^2 x^I x^A} > \frac{1}{8\bar{\alpha}\psi x^I}(\frac{1}{4\bar{\alpha}\psi x^I}-1) \tag{27}$$

Let $T = 4\bar{\alpha}\psi$. The above can be rearranged as

$$\frac{1-2x^A T(1-x^I T)-(x^I T)^2}{2x^A} > \frac{1-x^I T}{x^I} \tag{28}$$

Which reduces to

$$2x^A(1-(x^I T)^2) < x^I(1-(x^I T)^2) \tag{29}$$

Since $(x^I T)^2 = (4\bar{\alpha}\psi x^I)^2 < 1$, the above can never be satisfied when $x^A > x^I$.

Case 2: $(x_1^*|D=1) = x^I$ and $(x_1^*|D=0) = x^I + 8\bar{\alpha}\psi(x^I)^2\gamma$

The equilibrium condition for the faction is

$$\begin{aligned} -(x^I - x^A)^2 - (x^I + x^A)^2 > \\ -(x^I + 8\bar{\alpha}\psi(x^I)^2\gamma - x^A)^2 - (1-4\bar{\alpha}\psi\gamma(x^I + 8\bar{\alpha}\psi(x^I)^2\gamma))(x^I + x^A)^2 \\ -4\bar{\alpha}\psi\gamma(x^I + 8\bar{\alpha}\psi(x^I)^2\gamma)(x^I - x^A)^2 \end{aligned} \tag{30}$$

Denoting $\Delta = 8\bar{\alpha}\psi\gamma(x^I)^2$, the above can be rearranged as

$$0 > -\Delta^2 - 2\Delta(x^I - x^A) + 16\bar{\alpha}\psi\gamma(x^I + \Delta)x^I x^A \tag{31}$$

Substituting $\Delta = 8\bar{\alpha}\psi\gamma(x^I)^2$ and dividing by $16\bar{\alpha}\psi\gamma(x^I)^2$, the above reduces to

$$x^A < \frac{x^I}{2} \tag{32}$$

Given $x^A > x^I$, the condition can never be satisfied.

Thus, dissent never emerges in equilibrium if $x^A > x^I$ and $\gamma < \frac{1}{2}$.

Consider now the conditions on the cost of dissent δ. The proof of the first condition ($\delta > 4\bar{\alpha}x^I(2\gamma - 1)$) is presented in the main body of the paper. Here, I present a formal proof of the second condition. Let $\delta \geq 4\bar{\alpha}x^I$. After experiencing dissent, the incumbent would turn into a sure loser. Therefore:

$$(x_1^*|D = 1) = x^I \tag{33}$$

Conversely, (from Lemma 2) if the leading incumbent experiences no dissent:

$$(x_1^*|D = 0) = x^I - 8\bar{\alpha}\psi(x^I)^2(1 - \gamma) \tag{34}$$

Dissent strictly increases the ally's utility if and only if:

$$-(x^I - x^A)^2 - (x^I + x^A)^2 > \text{(35)}$$

$$-(x^I - 8\bar{\alpha}\psi(x^I)^2(1-\gamma) - x^A)^2 - (1 - 4\bar{\alpha}\psi(1-\gamma)(x^I - 8\bar{\alpha}\psi(x^I)^2(1-\gamma))(x^I - x^A)^2$$

$$-4\bar{\alpha}\psi(1-\gamma)(x^I - 8\bar{\alpha}\psi(x^I)^2(1-\gamma))(x^I + x^A)^2$$

This reduces to

$$x^A[1-2(4\bar{\alpha}\psi x^I(1-\gamma))(1-4\bar{\alpha}\psi x^I(1-\gamma))]+4\bar{\alpha}\psi(x^I)^2(1-\gamma)(1-4\bar{\alpha}\psi x^I(1-\gamma)) < 0 \tag{36}$$

The LHS is increasing in x^A and never satisfied at $x^A = 0$. Hence, dissent by an extremist ally emerges only if $\delta < 4\bar{\alpha}x^I$.

Finally, I must prove that there exist unique $\underline{\gamma}$, $\overline{\gamma}$, $\underline{x^I}$ and $\underline{x^A}$ such that dissent by an extremist ally emerges in equilibrium only if $\underline{\gamma} < \gamma < \overline{\gamma}$, $x^I > \underline{x^I}$ and $x^A > \underline{x^A}$.

Suppose that $4\bar{\alpha}x^I(2\gamma - 1) < \delta < 4\bar{\alpha}x^I$ and $\gamma > \frac{1}{2}$, and conjecture the existence of an equilibrium in which the ally chooses to dissent. We must consider two cases:

1. $(x_1^*|D=1) = \frac{1}{4\bar{\alpha}\psi}$ and $(x_1^*|D=0) = x^I - 8\bar{\alpha}\psi(1-\gamma)(x^I)^2$, when $x^I + 8\bar{\alpha}\psi(x^I)^2\gamma \geq \frac{1}{4\bar{\alpha}\psi}$

2. $(x_1^*|D=1) = x^I + 8\bar{\alpha}\psi(x^I)^2\gamma$ and $(x_1^*|D=0) = x^I - 8\bar{\alpha}\psi(1-\gamma)(x^I)^2$, when $x^I + 8\bar{\alpha}\psi(x^I)^2\gamma < \frac{1}{4\bar{\alpha}\psi}$

I will analyse each of the two cases separately.

Case 1: $(x_1^*|D=1) = \frac{1}{4\bar{\alpha}\psi}$ **and** $(x_1^*|D=0) = x^I - 8\bar{\alpha}\psi(1-\gamma)(x^I)^2$

From Lemma 2, the equilibrium conditions for the incumbent are

$$\gamma > \frac{1}{2} \tag{37}$$

$$\gamma > \frac{1}{8\bar{\alpha}\psi x^I}(\frac{1}{4\bar{\alpha}\psi x^I} - 1) \tag{38}$$

Additionally, the equilibrium condition for the faction is

$$-(\frac{1}{4\bar{\alpha}\psi} - x^A)^2 - \gamma(x^I - x^A)^2 - (1-\gamma)(x^I + x^A)^2 > \tag{39}$$

$$-(x^I - 8\bar{\alpha}\psi(x^I)^2(1-\gamma) - x^A)^2 - [1 - 4\bar{\alpha}\psi(1-\gamma)(x^I - 8\bar{\alpha}\psi(x^I)^2(1-\gamma))](x^I - x^A)^2$$

$$-4\bar{\alpha}\psi(1-\gamma)(x^I - 8\bar{\alpha}\psi(x^I)^2(1-\gamma))(x^I + x^A)^2$$

Let $I = 4\bar{\alpha}\psi(x^I - 8\bar{\alpha}\psi(x^I)^2(1-\gamma))$. We can rewrite the above condition as:

$$(1-\gamma)(1-I)((x^I - x^A)^2 - (x^I + x^A)^2) > (\frac{1}{4\bar{\alpha}\psi} - x^A)^2 - (\frac{I}{4\bar{\alpha}\psi} - x^A)^2 \tag{40}$$

Which is equivalent to

$$(1-\gamma)(1-I)(-4x^Ax^I) > \frac{-x^A}{2\bar{\alpha}\psi}(1-I) + \frac{1}{(4\bar{\alpha}\psi)^2}(1+I)(1-I) \tag{41}$$

By substituting $I = 4\bar{\alpha}\psi(x^I - 8\bar{\alpha}\psi(x^I)^2(1-\gamma))$ and solving for γ we get the following condition:

$$\gamma > \frac{1 + (2x^A - x^I)(8\bar{\alpha}\psi x^I - 1)(4\bar{\alpha}\psi)}{2x^I(4\bar{\alpha}\psi)^2(2x^A - x^I)} \tag{42}$$

Thus, the conjectured equilibrium exists if and only if the following conditions are satisfied:

1. $\gamma > \underline{\gamma} = max \quad \{\frac{1}{2}, \frac{1}{8\bar{\alpha}\psi x^I}(\frac{1}{4\bar{\alpha}\psi x^I} - 1), \frac{1+(2x^A-x^I)(8\bar{\alpha}\psi x^I-1)(4\bar{\alpha}\psi)}{2x^I(4\bar{\alpha}\psi)^2(2x^A-x^I)}\}$

2. $x^I > \frac{1}{8\bar{\alpha}\psi}$

3. $x^A > \frac{1}{8\bar{\alpha}\psi} + \frac{x^I}{2}$

Where the conditions on x^I and x^A ensure that the range $[\underline{\gamma}, 1]$ exists.

Case 2: $(x_1^*|D=1) = x^I + 8\bar{\alpha}\psi(x^I)^2\gamma$ **and** $(x_1^*|D=0) = x^I - 8\bar{\alpha}\psi(1-\gamma)(x^I)^2$

From Lemma 4, the equilibrium conditions for the incumbent are

$$\gamma > \frac{1}{2} \tag{43}$$

$$\gamma < \frac{1}{8\bar{\alpha}\psi x^I}(\frac{1}{4\bar{\alpha}\psi x^I} - 1) \tag{44}$$

Additionally, the the equilibrium condition for the faction is

$$\begin{aligned} & -(x^I + 8\bar{\alpha}\psi(x^I)^2\gamma - x^A)^2 \quad (45) \\ -(1 - 4\bar{\alpha}\psi(x^I + 8\alpha\psi(x^I)^2\gamma)\gamma)(x^I + x^A)^2 - 4\bar{\alpha}\psi(x^I + 8\bar{\alpha}\psi(x^I)^2\gamma)\gamma(x^I - x^A)^2 > \\ -(x^I - 8\bar{\alpha}\psi(x^I)^2(1-\gamma) - x^A)^2 - (1 - 4\bar{\alpha}\psi(x^I - 8\bar{\alpha}\psi(x^I)^2(1-\gamma))(1-\gamma))(x^I - x^A)^2 \\ -4\bar{\alpha}\psi(x^I - 8\bar{\alpha}\psi(x^I)^2(1-\gamma))(1-\gamma)(x^I + x^A)^2 \end{aligned}$$

Let $I = 4\bar{\alpha}\psi(x^I + 8\bar{\alpha}\psi(x^I)^2\gamma)$ and $x^D = x^I + 8\bar{\alpha}\psi(x^I)^2\gamma$. We can rewrite the above condition as:

$$\begin{aligned} -(x^D - x^A)^2 - (1 - \gamma I)(x^I + x^A)^2 - \gamma I(x^I - x^A)^2 > \quad (46) \\ -(x^D - x^A - 8\bar{\alpha}\psi(x^I)^2)^2 - (1 - (1-\gamma)(I - 4\bar{\alpha}\psi(8\bar{\alpha}\psi(x^I)^2)))(x^I - x^A)^2 \\ -(1-\gamma)(I - 4\bar{\alpha}\psi(8\bar{\alpha}\psi(x^I)^2))(x^I + x^A)^2 \end{aligned}$$

By expanding, letting $x^D = \frac{I}{4\bar{\alpha}\psi}$ and dividing both sides by $4x^I$ we get:

$$\gamma I x^A > x^A - x^A((1-\gamma)(I - 2(4\bar{\alpha}\psi x^I)^2) + x^I(I - 4\bar{\alpha}\psi x^A - (4\bar{\alpha}\psi x^I)^2) \quad (47)$$

Which is equivalent to:

$$x^A + I(x^I - x^A) - 4\bar{\alpha}\psi x^I x^A + (4\bar{\alpha}\psi x^I)^2(2x^A(1-\gamma) - x^I) < 0 \quad (48)$$

By substituting $I = 4\bar{\alpha}\psi(x^I + 8\bar{\alpha}\psi(x^I)^2\gamma)$ and solving for γ we get the following condition:

$$\gamma > \frac{4\bar{\alpha}\psi x^I(2x^A - x^I)(4\bar{\alpha}\psi x^I - 1) + x^A}{32\bar{\alpha}^2\psi^2(x^I)^2(2x^A - x^I)} \quad (49)$$

Thus, the conjectured equilibrium exists if and only if the following conditions are satisfied:

1. $\underline{\gamma} = max \quad \{\frac{1}{2}, \frac{4\bar{\alpha}\psi x^I(2x^A - x^I)(4\bar{\alpha}\psi x^I - 1) + x^A}{32\bar{\alpha}^2\psi^2(x^I)^2(2x^A - x^I)}\} < \gamma < min \quad \{1, \frac{1}{8\bar{\alpha}\psi x^I}(\frac{1}{4\bar{\alpha}\psi x^I} - 1)\} = \overline{\gamma}$

2. $\frac{\sqrt{3}-1}{8\bar{\alpha}\psi} < x^I < \frac{\sqrt{5}-1}{8\bar{\alpha}\psi}$

3. $x^A > \frac{x^I(4\bar{\alpha}\psi x^I)(4\bar{\alpha}\psi x^I+1)}{2(4\bar{\alpha}\psi x^I)(4\bar{\alpha}\psi x^I+1)-1}$

Where the conditions on x^I and x^A ensure that the range $[\underline{\gamma}, \overline{\gamma}]$ exists.

This conclude s the proof of Proposition 1. □

The following corollary also holds, with respect to Case 2:

Corollary 1A: $\underline{\gamma} = \frac{1}{2} \implies \overline{\gamma} < 1$

Proof. Corollary 1A tells us that it can never be the case that (i) $\underline{\gamma} = \frac{1}{2}$ and (ii) $\overline{\gamma} = 1$. For (i) to be true we need:

$$\frac{1}{2} > \frac{4\bar{\alpha}\psi x^I(2x^A - x^I)(4\bar{\alpha}\psi x^I - 1) + x^A}{32\bar{\alpha}^2\psi^2(x^I)^2(2x^A - x^I)} \tag{50}$$

Which reduces to:

$$x^A[1 - 8\bar{\alpha}\psi x^I] + 4\bar{\alpha}\psi(x^I)^2 < 0 \tag{51}$$

Which clearly requires:

$$x^I > \frac{1}{8\bar{\alpha}\psi} \tag{52}$$

For (ii) to be true we need:

$$1 < \frac{1}{8\bar{\alpha}\psi x^I}(\frac{1}{4\bar{\alpha}\psi x^I} - 1) \tag{53}$$

Which reduces to

$$x^I < \frac{1}{8\bar{\alpha}\psi} \tag{54}$$

Clearly the two conditions can never be simultaneously satisfied.

□

Proposition 2:

- *The likelihood of observing dissent (weakly) increases as the ally becomes more extreme*
- *There exists a unique $\widehat{x^I}(x^A) > \underline{x^I}$ such that if $x^I < \widehat{x^I}(x^A)$, then the likelihood of observing dissent increases as the incumbent becomes more extreme*

Proof. Denote $\Gamma(x^I, x^A, \bar{\alpha}\psi)$ the set of values of γ such that dissent is an equilibrium strategy iff $\gamma \in \Gamma$. From the proof of Proposition 1 is easy to verify that Γ is always weakly increasing in x^A. Analysing Cases 1 and 2, Γ is (weakly) increasing in x^I if and only if either one of the following sets of conditions is satisfied:

1. Case 1: $\underline{\gamma} = 1$. This requires $\frac{\sqrt{5}-1}{8\bar{\alpha}\psi} < x^I < \frac{1}{4\bar{\alpha}\psi}$ and $x^A > \frac{1+4\bar{\alpha}\psi x^I(1-4\bar{\alpha}\psi x^I}{8\bar{\alpha}\psi(1-4\bar{\alpha}\psi x^I)}$
2. Case 1: $\underline{\gamma} = \frac{1-4\bar{\alpha}\psi x^I}{2(4\bar{\alpha}\psi x^I)^2}$ which requires $x^I < \frac{\sqrt{5}-1}{8\bar{\alpha}\psi}$ and $x^A > \frac{x^I(1-(4\bar{\alpha}\psi x^I)^2)}{1-2(4\bar{\alpha}\psi x^I)^2}$. It is easy to verify that when $\underline{\gamma} = \frac{1-4\bar{\alpha}\psi x^I}{2(4\bar{\alpha}\psi x^I)^2}$ in Case 1, irrespective of the bounds in Case 2 Γ will be weakly increasing in x^I
3. Case 2: $\underline{\gamma} = \frac{4\bar{\alpha}\psi x^I(2x^A-x^I)(4\bar{\alpha}\psi x^I-1)+x^A}{32\bar{\alpha}^2\psi^2(x^I)^2(2x^A-x^I)}$ and $\overline{\gamma} = 1$, which requires $x^I < \frac{1}{8\bar{\alpha}\psi}$.

.

Thus, necessary and sufficient condition for Γ to be increasing in x^I is that $x^I < \widehat{x^I}(x^A)$.

$x^A > \frac{1+4\bar{\alpha}\psi x^I(1-4\bar{\alpha}\psi x^I}{8\bar{\alpha}\psi(1-4\bar{\alpha}\psi x^I)} \implies \widehat{x^I}(x^A) = \frac{1}{4\bar{\alpha}\psi}$, $\frac{x^I(1-(4\bar{\alpha}\psi x^I)^2}{1-2(4\bar{\alpha}\psi x^I)^2} < x^A < \frac{1+4\bar{\alpha}\psi x^I(1-4\bar{\alpha}\psi x^I}{8\bar{\alpha}\psi(1-4\bar{\alpha}\psi x^I)} \implies \widehat{x^I}(x^A) = \frac{\sqrt{5}-1}{8\bar{\alpha}\psi}$, and $x^A < \frac{x^I(1-(4\bar{\alpha}\psi x^I)^2}{1-2(4\bar{\alpha}\psi x^I)^2} \implies \widehat{x^I}(x^A) = \frac{1}{8\bar{\alpha}\psi}$ □

Proposition 3: *In equilibrium the voter benefits from the presence of an extreme ally to the incumbent party if:*

- *The cost of dissent δ is sufficiently large that it turns the leading incumbent into a trailing one, but not so large that it always hurts the voter ex ante ($\underline{\delta} < \delta < \overline{\delta_w}$)*
- *The value of information is sufficiently high*
  - *The prior (γ) is sufficiently close to $\frac{1}{2}$ ($\frac{1}{2} < \gamma < \overline{\gamma_w}$)*
  - *Incumbent and challenger are moderately polarized ($\underline{x^I_w} < x^I < \overline{x^I_w}$)*
  - *Learning the true state has a sufficiently large impact on the voter's preferences ($\bar{\alpha} > \bar{\alpha}_w$)*
- *The incumbent's ally is sufficiently extreme ($x^A > \underline{x^A_w}$)*

Proof. In order to identify *sufficient* conditions for the voter to benefit from dissent, suppose that $\gamma < \frac{1}{8\bar{\alpha}\psi x^I}(\frac{1}{4\bar{\alpha}\psi x^I} - 1)$. Then, $(x_1^*|D = 1) = x^I + 8\bar{\alpha}\psi(x^I)^2\gamma$ and $(x_1^*|D = 0) = x^I - 8\bar{\alpha}\psi(x^I)^2(1-\gamma)$. Dissent increases the voter's welfare if and only if the following condition is satisfied:

$$
\begin{aligned}
-4\bar{\alpha}\psi(x^I + 8\bar{\alpha}\psi(x^I)^2\gamma)\gamma\delta - \gamma(x^I + 8\bar{\alpha}\psi(x^I)^2\gamma - \bar{\alpha})^2 \\
-(1-\gamma)(x^I + 8\bar{\alpha}\psi(x^I)^2\gamma + \bar{\alpha})^2 - 4\bar{\alpha}\psi(x^I + 8\bar{\alpha}\psi(x^I)^2\gamma)(x^I - \bar{\alpha})^2 \\
-(1 - 4\bar{\alpha}\psi(x^I + 8\bar{\alpha}\psi(x^I)^2\gamma))(\gamma(x^I + \bar{\alpha})^2 + (1-\gamma)(x^I - \bar{\alpha})^2) > \\
-\gamma(x^I - 8\bar{\alpha}\psi(x^I)^2(1-\gamma) - \bar{\alpha})^2 - (1-\gamma)(x^I - 8\bar{\alpha}\psi(x^I)^2(1-\gamma) + \bar{\alpha})^2 \\
-4\bar{\alpha}\psi(x^I - 8\bar{\alpha}\psi(x^I)^2(1-\gamma))(x^I - \bar{\alpha})^2 \\
-(1 - 4\bar{\alpha}\psi(x^I - 8\bar{\alpha}\psi(x^I)^2(1-\gamma)))(\gamma(x^I - \bar{\alpha})^2 + (1-\gamma)(x^I + \bar{\alpha})^2)
\end{aligned}
\tag{55}
$$

Which reduces to:

$$
\delta < \frac{(1-2\gamma)(1 - 8\bar{\alpha}\psi x^I + 2(4\bar{\alpha}\psi x^I)^2 + 16\bar{\alpha}\psi^2(x^I)^3) - 4\psi(x^I)^2 + 4(4\bar{\alpha}\psi x^I\gamma)^2}{\psi\gamma(1 + 8\bar{\alpha}\psi x^I\gamma)} = \overline{\delta_w} \tag{56}
$$

If the above is satisfied, the voter benefits from dissent. However, we need to make sure that dissent would indeed emerge in equilibrium (given the incumbent's equilibrium policy choices with and without dissent). From the proof of Proposition 1 we know that this requires the following conditions:

1. $4\bar{\alpha}x^I(2\gamma - 1) \leq \delta < 4\bar{\alpha}x^I$

2. $\underline{\gamma} = max \quad \{\frac{1}{2}, \frac{4\bar{\alpha}\psi x^I(2x^A - x^I)(4\bar{\alpha}\psi x^I - 1) + x^A}{32\bar{\alpha}^2\psi^2(x^I)^2(2x^A - x^I)}\} < \gamma < min \quad \{1, \frac{1}{8\bar{\alpha}\psi x^I}(\frac{1}{4\bar{\alpha}\psi x^I} - 1)\} = \overline{\gamma}$

3. $\frac{\sqrt{3}-1}{8\bar{\alpha}\psi} < x^I < \frac{\sqrt{5}-1}{8\bar{\alpha}\psi}$

4. $x^A > \frac{x^I(4\bar{\alpha}\psi x^I)(4\bar{\alpha}\psi x^I + 1)}{2(4\bar{\alpha}\psi x^I)(4\bar{\alpha}\psi x^I + 1) - 1}$

Thus, for the voter to benefit the presence of the incumbent's extreme ally, both (46) and conditions 1 to 4 above must be satisfied. This requires $\overline{\delta_w} > 4\bar{\alpha}x^I(2\gamma - 1)$,

which reduces to:

$$(1-2\gamma)(1+4\bar{\alpha}\psi x^I(8\bar{\alpha}\psi x^I-2+4\psi(x^I)^2+8\bar{\alpha}\psi x^I\gamma^2+\gamma))-4\psi(x^I)^2+4(4\bar{\alpha}\psi x^I\gamma)^2>0 \quad (57)$$

The LHS is decreasing in γ, therefore the above establishes an upper bound $\overline{\gamma_w}$. For the condition to be possible to satisfy in equilibrium we need $\overline{\gamma_w}>\underline{\gamma}$. From the proof of Case 2 we can verify that $\underline{\gamma}=\frac{1}{2}$ when $x^I>\frac{1}{8\bar{\alpha}\psi}$ and $x^A>\frac{4\bar{\alpha}\psi(x^I)^2}{8\bar{\alpha}\psi x^I-1}$. Additionally, given Corollary 1A $\underline{\gamma}=\frac{1}{2}\implies\overline{\gamma}=\frac{1}{8\bar{\alpha}\psi x^I}(\frac{1}{4\bar{\alpha}\psi x^I}-1)$. Thus, the voter benefits from the presence of the incumbent's extreme ally if:

1. $4\bar{\alpha}x^I(2\gamma-1)\leq\delta<\overline{\delta_w}$
2. $\gamma<\overline{\gamma_w}$
3. $\underline{\gamma}=\frac{1}{2}<\gamma<\frac{1}{8\bar{\alpha}\psi x^I}(\frac{1}{4\bar{\alpha}\psi x^I}-1)=\overline{\gamma}$
4. $\frac{1}{8\bar{\alpha}\psi x^I}(\frac{1}{4\bar{\alpha}\psi x^I}-1)>\frac{1}{2}$
5. $\overline{\gamma_w}>\frac{1}{2}$
6. $x^I>\frac{1}{8\bar{\alpha}\psi}=\underline{x^I_w}$
7. $x^A>\frac{4\bar{\alpha}\psi(x^I)^2}{8\bar{\alpha}\psi x^I-1}=\underline{x^A_w}$

$\overline{\gamma_w}>\frac{1}{2}$ if and only if the following is satisfied:

$$-4\psi(x^I)^2+4(2\bar{\alpha}\psi x^I)^2>0 \quad (58)$$

Which reduces to

$$\alpha>\frac{1}{2\sqrt{\psi}}=\bar{\alpha}_w \quad (59)$$

$\frac{1}{8\bar{\alpha}\psi x^I}(\frac{1}{4\bar{\alpha}\psi x^I} - 1) > \frac{1}{2}$ if and only if

$$x^I < \frac{\sqrt{5}-1}{8\bar{\alpha}\psi} = \overline{x^I_w} \quad (60)$$

Thus we can rewrite the sufficient conditions for the voter to benefit from dissent in equilibrium as:

1. $\underline{\delta} < \delta < \overline{\delta_w}$
2. $\frac{1}{2} < \gamma < min\{\overline{\gamma}, \overline{\gamma_w}\}$
3. $\underline{x^I_w} < x^I < \overline{x^I_w}$
4. $x^A > \underline{x^A_w}$
5. $\bar{\alpha} > \bar{\alpha}_w$

□

Extension: What if the Ally Has Bargaining Power?

In this section I will consider the case in which the incumbent's ally has bargaining power over the first period policy making. I will thus assume that in the first period the incumbent maximises a weighted average of his own and the ally's utility:

$$U_1^W = \beta[-(x_1 - x^A)^2 + U_2^A(x_1, x_2, x^A)] + (1-\beta)[-(x_1 - x^I)^2 + U_2^I(x_1, x_2, x^I)] \quad (61)$$

This is equivalent to analysing a game in which, after the ally chooses wether to dissent, it engages in a bargaining stage with the incumbent to determine the policy to be implemented in the first period. Therefore, the parameter β represents, in this reduced form, the ally's bargaining power in the first period. As in the baseline model, I assume $x^C = -x^I \leq 0$ and $x^I < \frac{1}{4\alpha\psi}$. Additionally, I assume that in the second period the ally has no bargaining power. A discussion of the necessity and significance of this assumption is in the main body of the paper.

We can determine the equilibrium policy choice of the incumbent, proceeding as in the proof of Lemma 3.

Consider first a trailing incumbent. The following holds:

- Let $\beta x^A + (1-\beta)x^I \geq \frac{1}{4\bar{\alpha}a\psi}$, then $x_1^* = \beta x^A + (1-\beta)x^I$
- Let $\beta x^A + (1-\beta)x^I < \frac{1}{4\bar{\alpha}\psi}$, then $x_1^* = \min \{\frac{1}{4\bar{\alpha}\psi}; [\beta x^A + (1-\beta)x^I][1 + 8\bar{\alpha}\psi x^I \gamma]\}$

Consider now a leading incumbent:

- Let $\beta x^A + (1-\beta)x^I \geq \frac{1}{4\bar{\alpha}\psi}$.
  Then $x_1^* = \beta x^A + (1-\beta)x^I$ if $\gamma > \frac{1+4\bar{\alpha}\psi[(\beta x^A+(1-\beta)x^I)(4\alpha\psi x^I-1)]}{(4\bar{\alpha}\psi)^2[x^I(\beta x^A+(1-\beta)x^I)]}$, and $x_1^* = [\beta x^A + (1-\beta)x^I][1 - 8\bar{\alpha}\psi x^I(1-\gamma)]$ otherwise[10]
- Let $\beta x^A + (1-\beta)x^I < \frac{1}{4\bar{\alpha}\psi}$, then $x_1^* = [\beta x^A + (1-\beta)x^I][1 - 8\alpha\psi x^I(1-\gamma)]$

Proposition 4: *For all $x^I \geq 0$, there exist non-measure zero sets $\Gamma(x^I)$ and $B(x^I)$ such that if $\gamma \in \Gamma(x^I)$ and $\beta \in B(x^I)$ then dissent by an extreme ally occurs in equilibrium*

Proof. I proceed as in the proof of Proposition 1. Suppose that $4\bar{\alpha}x^I(2\gamma-1) < \delta < 4\bar{\alpha}x^I$ and $\gamma > \frac{1}{2}$, and conjecture the existence of an equilibrium in which the ally chooses to dissent. We must consider three cases:

1. $(x_1^*|D=1) = \beta x^A + (1-\beta)x^I$ and $(x_1^*|D=0) = [\beta x^A + (1-\beta)x^I][1 - 8\bar{\bar{\alpha}}\psi(1-\gamma)x^I]$

[10]When $\beta x^A + (1-\beta)x^I \geq \frac{1}{4\bar{\alpha}\psi}$ the leading incumbent's overall utility as a function of the first period policy has two maxima: one at $\beta x^A + (1-\beta)x^I$ and a second at $[\beta x^A + (1-\beta)x^I][1 - 8\bar{\alpha}\psi x^I(1-\gamma)]$. The condition on γ identifies which one of the two is the global maximum.

2. $(x_1^*|D=1) = \frac{1}{4\bar{\alpha}\psi}$ and $(x_1^*|D=0) = [\beta x^A + (1-\beta)x^I][1 - 8\bar{\bar{\alpha}}\psi(1-\gamma)x^I]$

3. $(x_1^*|D=1) = [\beta x^A + (1-\beta)x^I][1 + 8\bar{\alpha}\psi\gamma x^I]$ and $(x_1^*|D=0) = [\beta x^A + (1-\beta)x^I][1 - 8\bar{\alpha}\psi(1-\gamma)x^I]$

I will analyse each of the three cases separately.

Case 1: $(x_1^*|D=1) = \beta x^A + (1-\beta)x^I$, $(x_1^*|D=0) = (\beta x^A + (1-\beta)x^I)(1 - 8\alpha\psi x^I(1-\gamma))$

The equilibrium conditions for the incumbent are:

$$\gamma > \frac{1}{2} \tag{62}$$

$$\beta \geq \frac{1 - 4\bar{\alpha}\psi x^I}{4\alpha\psi(x^A - x^I)} \tag{63}$$

$$\gamma < \frac{1 + 4\bar{\alpha}\psi((\beta x^A + (1-\beta)x^I)(4\bar{\alpha}\psi x^I - 1))}{(4\bar{\alpha}\psi)^2 x^I(\beta x^A + (1-\beta)x^I)} \tag{64}$$

Additionally, the equilibrium condition for the faction is

$$\begin{aligned} &-(\beta x^A + (1-\beta)x^I - x^A)^2 - \gamma(x^I - x^A)^2 \\ -(1-\gamma)(x^I + x^A)^2 > &-[(\beta x^A + (1-\beta)x^I)(1 - 8\bar{\alpha}\psi x^I(1-\gamma)) - x^A]^2 \\ &-[1 - 4\bar{\alpha}\psi(1-\gamma)(\beta x^A + (1-\beta)x^I)(1 - 8\bar{\alpha}\psi x^I(1-\gamma))](x^I - x^A)^2 \\ &-[4\bar{\alpha}\psi(1-\gamma)(\beta x^A + (1-\beta)x^I)(1 - 8\bar{\alpha}\psi x^I(1-\gamma))](x^I + x^A)^2 \end{aligned} \tag{65}$$

Let $x^D = \beta x^A + (1-\beta)x^I$ and $x^D - \Delta = (\beta x^A + (1-\beta)x^I)(1 - 8\bar{\alpha}\psi x^I(1-\gamma))$ where $\Delta = (\beta x^A + (1-\beta)x^I)8\bar{\alpha}\psi x^I(1-\gamma)$. The above reduces to

$$-\Delta^2 + 2\Delta(x^D - x^A) + 4x^I x^A(1-\gamma) - 16\bar{\alpha}\psi(1-\gamma)x^I x^A(x^D - \Delta) < 0 \tag{66}$$

Substituting $\Delta = (\beta x^A + (1-\beta)x^I)8\bar{\alpha}\psi x^I(1-\gamma)$ and dividing for $4x^I(1-\gamma)$ gives

$$-x^I(4\bar{\alpha}\psi)^2(1-\gamma)(\beta x^A + (1-\beta)x^I)^2 + 4\bar{\alpha}\psi(\beta x^A + (1-\beta)x^I)(x^D - x^A) \qquad (67)$$
$$+x^A - 4\bar{\alpha}\psi x^A(x^D - (\beta x^A + (1-\beta)x^I)8\bar{\alpha}\psi x^I(1-\gamma)) < 0$$

Substituting $x^D = \beta x^A + (1-\beta)x^I$ and solving for γ gives us condition:

$$\gamma > 1 + \frac{x^A - 4\bar{\alpha}\psi[\beta x^A + (1-\beta)x^I][2x^A - \beta x^A - (1-\beta)x^I]}{(4\bar{\alpha}\psi)^2 x^I[\beta x^A + (1-\beta)x^I][2x^A - \beta x^A - (1-\beta)x^I]} \qquad (68)$$

Thus, the conjectured equilibrium exist if and only if the following conditions are satisfied:

1. $\underline{\gamma} = max \quad \{\frac{1}{2}, 1 + \frac{x^A - 4\bar{\alpha}\psi[\beta x^A + (1-\beta)x^I][2x^A - \beta x^A - (1-\beta)x^I]}{(4\bar{\alpha}\psi)^2 x^I[\beta x^A + (1-\beta)x^I][2x^A - \beta x^A - (1-\beta)x^I]}\} < \gamma < \frac{1+4\bar{\alpha}\psi((\beta x^A + (1-\beta)x^I)(4\bar{\alpha}\psi x^I - 1))}{(4\bar{\alpha}\psi)^2 x^I(\beta x^A + (1-\beta)x^I)} = \overline{\gamma}$
2. $\underline{\beta} = \frac{1-4\bar{\alpha}\psi x^I}{4\bar{\alpha}\psi(x^A - x^I)} \leq \beta < min \quad \{1, \frac{1+4\bar{\alpha}\psi x^I(2\bar{\alpha}\psi x^I - 1)}{4\bar{\alpha}\psi(x^A - x^I)(1-2\bar{\alpha}\psi x^I)}\} = \overline{\beta}$
3. $x^A > \frac{1}{4\bar{\alpha}\psi}$

The conditions on β ensure that the range $[\underline{\gamma}, \overline{\gamma}]$ exists. The condition on x^A ensures that the range $[\underline{\beta}, \overline{\beta}]$ exists.

Case 2: $(x_1^*|D=1) = \frac{1}{4\bar{\alpha}\psi}$, $(x_1^*|D=0) = (\beta x^A + (1-\beta)x^I)(1 - 8\bar{\alpha}\psi x^I(1-\gamma)$

The equilibrium conditions for the incumbent are:

$$\gamma > \frac{1}{2} \qquad (69)$$
$$\beta < \frac{1-4\bar{\alpha}\psi x^I}{4\bar{\alpha}\psi(x^A - x^I)} \qquad (70)$$
$$\gamma > \frac{1}{8\bar{\alpha}\psi x^I}(\frac{1}{4\bar{\alpha}\psi(\beta x^A + (1-\beta)x^I)} - 1) \qquad (71)$$

Additionally, the equilibrium condition for the faction is

$$-(\frac{1}{4\bar{\alpha}\psi} - x^A)^2 - \gamma(x^I - x^A)^2 - (1-\gamma)(x^I + x^A)^2 > \tag{72}$$
$$-[(\beta x^A + (1-\beta)x^I)(1 - 8\bar{\alpha}\psi x^I(1-\gamma)) - x^A]^2$$
$$-[1 - 4\bar{\alpha}\psi(1-\gamma)(\beta x^A + (1-\beta)x^I)(1 - 8\bar{\alpha}\psi x^I(1-\gamma))](x^I - x^A)^2$$
$$-[4\bar{\alpha}\psi(1-\gamma)(\beta x^A + (1-\beta)x^I)(1 - 8\bar{\alpha}\psi x^I(1-\gamma))](x^I + x^A)^2$$

Let $I = 4\bar{\alpha}\psi(\beta x^A + (1-\beta)x^I)(1 - 8\bar{\alpha}\psi x^I(1-\gamma))$. The above can be rewritten as:

$$-(\frac{1}{4\bar{\alpha}\psi} - x^A)^2 - \gamma(x^I - x^A)^2 - (1-\gamma)(x^I + x^A)^2 > \tag{73}$$
$$-(\frac{I}{4\bar{\alpha}\psi} - x^A)^2 - (1 - I(1-\gamma))(x^I - x^A)^2 - I(1-\gamma)(x^I + x^A)^2$$

Which reduces to

$$(1-I)(\frac{x^A}{2\bar{\alpha}\psi} - 4x^I x^A(1-\gamma) - \frac{1+I}{(4\bar{\alpha}\psi)^2}) > 0 \tag{74}$$

By substituting $I = 4\bar{\alpha}\psi(\beta x^A + (1-\beta)x^I)(1 - 8\bar{\alpha}\psi x^I(1-\gamma))$ and solving for γ we get condition:

$$1 + \frac{-1 + 4\bar{\alpha}\psi(2x^A - x^I - \beta(x^A - x^I)}{-2(4\bar{\alpha}\psi)^2 x^I(2x^A - x^I - \beta(x^A - x^I)} < \gamma < 1 \tag{75}$$

Thus, the conjectured equilibrium exists if and only if the following conditions are satisfied:

1. $\underline{\gamma} = max \quad \{\frac{1}{2}, 1 + \frac{-1+4\bar{\alpha}\psi(2x^A - x^I - \beta(x^A - x^I)}{-2(4\bar{\alpha}\psi)^2 x^I(2x^A - x^I - \beta(x^A - x^I)}, \frac{1}{8\bar{\alpha}\psi x^I}(\frac{1}{4\bar{\alpha}\psi(\beta x^A + (1-\beta)x^I)} - 1)\} < \gamma < 1 = \overline{\gamma}$

2. $\underline{\beta} = \max \quad \{0, \frac{1 - 4\bar{\alpha}\psi x^I - 2(4\bar{\alpha}\psi x^I)^2}{4\bar{\alpha}\psi(x^A - x^I)(8\bar{\alpha}\psi x^I + 1)}\} < \beta < \overline{\beta} = min \quad \{\frac{1 - 4\bar{\alpha}\psi x^I}{4\bar{\alpha}\psi(x^A - x^I)}, \frac{4\bar{\alpha}\psi(2x^A - x^I) - 1}{4\bar{\alpha}\psi(x^A - x^I)}\}$

3. $x^A > \underline{x^A} = max \quad \{\frac{1 + 4\bar{\alpha}\psi x^I}{8\bar{\alpha}\psi}, \frac{1 + 4\bar{\alpha}\psi x^I}{4\bar{\alpha}\psi(1 + 8\bar{\alpha}\psi x^I)}\}$

The conditions on β ensure that the range $[\underline{\gamma}, \overline{\gamma}]$ exists. The condition on x^A ensures that the range $[\underline{\beta}, \overline{\beta}]$ exists.

Case 3: $(x_1^*|D=1) = (\beta x^A + (1-\beta)x^I)(1 + 8\bar{\alpha}\psi x^I \gamma)$**,** $(x_1^*|D=0) = (\beta x^A + (1-\beta)x^I)(1 - 8\bar{\alpha}\psi x^I(1-\gamma)$

The equilibrium conditions for the incumbent are:

$$\gamma > \frac{1}{2} \tag{76}$$

$$\beta < \frac{1 - 4\bar{\alpha}\psi x^I}{4\bar{\alpha}\psi(x^A - x^I)} \tag{77}$$

$$\gamma < \frac{1}{8\bar{\alpha}\psi x^I}\left(\frac{1}{4\bar{\alpha}\psi(\beta x^A + (1-\beta)x^I)} - 1\right) \tag{78}$$

Additionally, the equilibrium condition for the faction is

$$\begin{aligned} &-[(\beta x^A + (1-\beta)x^I)(1 + 8\bar{\alpha}\psi x^I \gamma)) \\ &-x^A]^2 - [1 - 4\bar{\alpha}\psi\gamma(\beta x^A + (1-\beta)x^I)(1 + 8\bar{\alpha}\psi x^I\gamma)](x^I + x^A)^2 \\ &-[4\alpha\psi\gamma(\beta x^A + (1-\beta)x^I)(1 + 8\bar{\alpha}\psi x^I\gamma)](x^I - x^A)^2 > \\ &-[(\beta x^A + (1-\beta)x^I)(1 - 8\bar{\alpha}\psi x^I(1-\gamma)) - x^A]^2 \\ &-[1 - 4\bar{\alpha}\psi(1-\gamma)(\beta x^A + (1-\beta)x^I)(1 - 8\bar{\alpha}\psi x^I(1-\gamma))](x^I - x^A)^2 \\ &-[4\bar{\alpha}\psi(1-\gamma)(\beta x^A + (1-\beta)x^I)(1 - 8\bar{\alpha}\psi x^I(1-\gamma))](x^I + x^A)^2 \end{aligned} \tag{79}$$

Let $x^D = (\beta x^A + (1-\beta)x^I)(1 + 8\bar{\alpha}\psi x^I\gamma)$. We can rewrite the above as:

$$\begin{aligned} &-(x^D - x^A)^2 - (1 - 4\bar{\alpha}\psi x^D\gamma)(x^I + x^A)^2 - 4\bar{\alpha}\psi x^D\gamma)(x^I - x^A)^2 > \\ &-(x^D - 8\bar{\alpha}\psi x^I(\beta x^A + (1-\beta)x^I) - x^A)^2 \\ &-(1 - 4\bar{\alpha}\psi(1-\gamma)(x^D - 8\bar{\alpha}\psi x^I(\beta x^A + (1-\beta)x^I))(x^I - x^A)^2 \\ &-4\bar{\alpha}\psi(1-\gamma)(x^D - 8\bar{\alpha}\psi x^I(\beta x^A + (1-\beta)x^I)(x^I + x^A)^2 \end{aligned} \tag{80}$$

Which reduces to

$$-4x^Ix^A + 16\bar{\alpha}\psi x^Dx^Ix^A\gamma > \tag{81}$$
$$-(8\bar{\alpha}\psi x^I(\beta x^A + (1-\beta)x^I))^2 + 16\bar{\alpha}\psi x^I(\beta x^A + (1-\beta)x^I)(x^D - x^A)$$
$$-16\bar{\alpha}\psi x^Ix^A(1-\gamma)(x^D - 8\bar{\alpha}\psi x^I(\beta x^A + (1-\beta)x^I))$$

By substituting $x^D = (\beta x^A + (1-\beta)x^I)(1 + 8\bar{\alpha}\psi(x^I)^2\gamma)$ and solving for γ we obtain condition:

$$\gamma > \frac{x^A + 4\bar{\alpha}\psi(\beta x^A + (1-\beta)x^I)(1-4\bar{\alpha}\psi x^I)(\beta x^A + (1-\beta)x^I - 2x^A)}{2x^I(4\bar{\alpha}\psi)^2(\beta x^A + (1-\beta)x^I)(-\beta x^A - (1-\beta)x^I + 2x^A)} \tag{82}$$

Thus the conjectured equilibrium exists if and only if the following conditions are satisfied:

1. $\underline{\gamma} = max \in \{\frac{1}{2}, \frac{x^A+4\bar{\alpha}\psi(\beta x^A+(1-\beta)x^I)(1-4\bar{\alpha}\psi x^I)(\beta x^A+(1-\beta)x^I-2x^A)}{2x^I(4\bar{\alpha}\psi)^2(\beta x^A+(1-\beta)x^I)(-\beta x^A-(1-\beta)x^I+2x^A)}\} < \gamma <$ $min \quad \{1, \frac{1}{8\bar{\alpha}\psi x^I}(\frac{1}{4\bar{\alpha}\psi(\beta x^A+(1-\beta)x^I)} - 1)\} = \overline{\gamma}$
2. $\underline{\beta} = max \in \{0, 1 - \frac{1}{2}\sqrt{\frac{x^I(4\bar{\alpha}\psi x^A)^2+4\bar{\alpha}\psi(x^A)^2-x^A}{\bar{\alpha}\psi(x^A-x^I)^2(1+4\bar{\alpha}\psi x^I)}}\} < \beta <$ $min \in \{\frac{1+2x^I(4\bar{\alpha}\psi)^2(x^A-x^I)-\sqrt{1+4(4x^Ax^I\bar{\alpha}\psi)^2(4\bar{\alpha}\psi)^2}}{32\bar{\alpha}^2\psi^2x^I(x^A-x^I)}, \frac{1-4\bar{\alpha}\psi x^I(1+4\bar{\alpha}\psi x^I)}{4\bar{\alpha}\psi(x^A-x^I)(1+4\bar{\alpha}\psi x^I)}\} = \overline{\beta}$
3. $x^A > max \in \{\frac{1}{4\bar{\alpha}\psi(1+4\bar{\alpha}\psi x^I)}, \frac{x^I(1-(4\bar{\alpha}\psi x^I)^2)}{1-2(4\bar{\alpha}\psi x^I)^2}, \frac{1+4\bar{\alpha}\psi x^I}{4\bar{\alpha}\psi(1+8\bar{\alpha}\psi x^I)}\}$
4. $x^I < \frac{\sqrt{5}-1}{8\bar{\alpha}\psi}$

The conditions on β ensure that the range $[\underline{\gamma}, \overline{\gamma}]$ exists. The conditions on x^A and x^I ensure that the range $[\underline{\beta}, \overline{\beta}]$ exists. □

Corollary 1: *Suppose that* $\frac{1}{8\bar{\alpha}\psi} < x^I$ *and* $\frac{1}{4\bar{\alpha}\psi} < x^A < \frac{1}{4\bar{\alpha}\psi(1-2\bar{\alpha}\psi x^I)}$. *Then, for all* $\beta \in [0, 1)$, *there exists a non-measure zero set* $\Gamma(\beta)$ *such that if* $\gamma \in \Gamma(\beta)$, *then dissent by an extreme ally occurs in equilibrium*

Proof. From an analysis of the cases above we can verify that sufficient conditions for the claim (for all $\beta \in [0, 1)$, there exists a non-measure zero set $\Gamma(\beta)$) to hold are:

- The binding upper bound $\overline{\beta}$ in case 1 is $= 1$
- The binding lower bound $\underline{\beta}$ in case 2 is $= 0$
- The binding upper bound $\overline{\beta}$ in case 2 is $= \frac{1-4\bar{\alpha}\psi x^I}{4\bar{\alpha}\psi(x^A - x^I)}$ (which is also the lower bound from case 1)

For the three conditions to be satisfied we need:

- $\frac{1}{4\alpha\psi} < x^A < \frac{1}{4\alpha\psi(1-2\alpha\psi x^I)}$
- $x^I > \frac{1}{8\alpha\psi}$

□

Appendix B: Dissent by a Moderate Ally

In this section I consider an ally whose bliss point is to the left of the incumbent: $0 < x^A < x^I$. In line with the rest of the paper, I maintain the assumption that $x^I < \frac{1}{4\bar{\alpha}\psi}$.

Proposition 1A: *There exist $\underline{\gamma}_m$, $\overline{\gamma}_m$, $\overline{x^A}_m$ and $\overline{x^I}_m$ such that the incumbent's moderate ally chooses to dissent in equilibrium if and only if:*

1. *The party is trailing, but its disadvantage is not too large ($\underline{\gamma}_m < \gamma < \overline{\gamma}_m$, where $\overline{\gamma}_m \leq \frac{1}{2}$)*
2. *Electoral cost of dissent sufficiently large to turn trailing incumbent into sure loser ($\delta \geq \overline{\delta}$)*

3. *Both the incumbent and his ally are sufficiently moderate* $(x^I < \overline{x^I}_m$ *and* $x^A < \overline{x^A}_m)$

Proof. The proof of the first point (incumbent must be trailing) is omitted, since it is obtained by applying the same logic used in proving Proposition 1 (i.e. it is easy to verify given the calculations in the proof of Proposition 1 that dissent never emerges if $x^A < x^I$ and $\gamma > \frac{1}{2}$). To prove the remainder of the proposition I must analyse all possible pairs of equilibrium policies. From above we know that in any equilibrium in which the ally chooses to dissent $(x_1^*|D = 1) = x^I$, and that dissent never occurs in equilibrium if $x^I \geq \frac{1}{4\bar{\alpha}\psi}$. Therefore, we must consider two cases:

- $(x_1^*|D = 1) = x^I$ and $(x_1^*|D = 0) = \frac{1}{4\bar{\alpha}\psi}$, when $x^I + 8\bar{\alpha}\psi(x^I)^2\gamma \geq \frac{1}{4\bar{\alpha}\psi}$
- $(x_1^*|D = 1) = x^I$ and $(x_1^*|D = 0) = x^I + 8\bar{\alpha}\psi(x^I)^2\gamma$, when $x^I + 8\bar{\alpha}\psi(x^I)^2\gamma < \frac{1}{4\bar{\alpha}\psi}$

I will analyse each case separately, conjecturing the existence of an equilibrium in which the ally chooses to dissent.

Case 1: $(x_1^*|D = 1) = x^I$ and $(x_1^*|D = 0) = \frac{1}{4\bar{\alpha}\psi}$

The equilibrium conditions for the incumbent are

$$\gamma \leq \frac{1}{2} \tag{83}$$

$$\gamma > \frac{1}{8\bar{\alpha}\psi x^I}(\frac{1}{4\bar{\alpha}\psi x^I} - 1) \tag{84}$$

Additionally, the equilibrium condition for the faction is

$$-(x^I - x^A)^2 - (x^I + x^A)^2 > -(\frac{1}{4\bar{\alpha}\psi} - x^A)^2 - \gamma(x^I - x^A)^2 - (1-\gamma)(x^I + x^A)^2 \tag{85}$$

Which reduces to

$$\gamma < \frac{-8\bar{\alpha}\psi x^A(1-4\bar{\alpha}\psi x^I) - (4\bar{\alpha}\psi x^I)^2 + 1}{(8\bar{\alpha}\psi)^2 x^I x^A} \tag{86}$$

Thus, the conjectured equilibrium exists if and only if the following conditions are satisfied:

1. $\frac{1}{8\bar{\alpha}\psi x^I}(\frac{1}{4\bar{\alpha}\psi x^I} - 1) < \gamma < min \quad \{\frac{1}{2}, \frac{-8\bar{\alpha}\psi x^A(1-4\bar{\alpha}\psi x^I)-(4\bar{\alpha}\psi x^I)^2+1}{(8\bar{\alpha}\psi)^2 x^I x^A}\}$
2. $\frac{\sqrt{5}-1}{8\bar{\alpha}\psi} < x^I < \frac{1}{4\bar{\alpha}\psi}$
3. $x^A < \frac{x^I}{2}$

The conditions on x^I and x^A ensure that the range $[\underline{\gamma}, \overline{\gamma}]$ exists.

Case 2: $(x_1^*|D=1) = x^I$ and $(x_1^*|D=0) = x^I + 8\bar{\alpha}\psi(x^I)^2\gamma$

The equilibrium conditions for the incumbent are

$$\gamma \leq \frac{1}{2} \tag{87}$$

$$\gamma < \frac{1}{8\bar{\alpha}\psi x^I}(\frac{1}{4\bar{\alpha}\psi x^I} - 1) \tag{88}$$

$$\tag{89}$$

Which requires

$$x^I < \frac{1}{4\bar{\alpha}\psi} \tag{90}$$

Additionally, the equilibrium condition for the faction is

$$
\begin{aligned}
-(x^I - x^A)^2 - (x^I + x^A)^2 > \\
-(x^I + 8\bar{\alpha}\psi(x^I)^2\gamma - x^A)^2 - (1 - 4\bar{\alpha}\psi\gamma(x^I + 8\bar{\alpha}\psi(x^I)^2\gamma))(x^I + x^A)^2 \\
-4\bar{\alpha}\psi\gamma(x^I + 8\bar{\alpha}\psi(x^I)^2\gamma)(x^I - x^A)^2
\end{aligned}
\tag{91}
$$

Which reduces to

$$
x^A < \frac{x^I}{2} \tag{92}
$$

Thus, the conjectured equilibrium exists if and only if the following conditions are satisfied:

1. $0 < \gamma < min \quad \{\frac{1}{2}, \frac{1}{8\bar{\alpha}\psi x^I}(\frac{1}{4\bar{\alpha}\psi x^I} - 1)\}$
2. $x^A < \frac{x^I}{2}$
3. $x^I < \frac{1}{4\bar{\alpha}\psi}$

This concludes the proof. □

Ideology For The Future

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, Wildawsky 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 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.

[1] Kartik, 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.

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: $x_V \in \{\underline{\alpha}, \overline{\alpha}\}$ where $\overline{\alpha} = -\underline{\alpha} \geq 0$. 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:

$$U_t^V = -(x_V - x_t)^2 + e_t \tag{1}$$

Where

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

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):

$$U_t^i = -(x_i - x_t)^2 \tag{2}$$
$$\forall i \in \{L, R\}$$

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 $x_V \in \{\underline{\alpha}, \overline{\alpha}\}$, (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: $x_L \leq \underline{\alpha} \leq 0 \leq \overline{\alpha} \leq x_R$.

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 $x_V = \overline{\alpha}$, 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

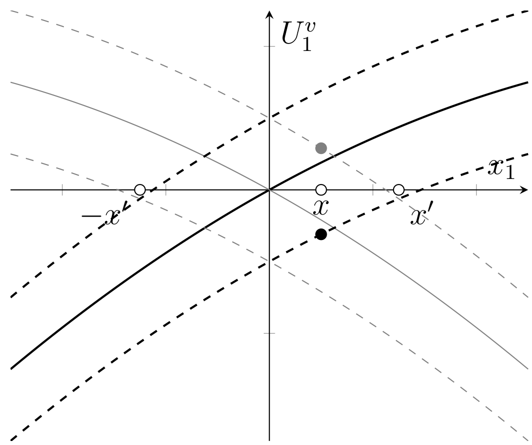

Figure 1: Voter's payoff realization as a function of first-period policy. The thick (thin) curves represent the case in which $x_V = \overline{\alpha}$ ($x_V = \underline{\alpha}$). Solid curves are the voter's expected payoff $E[U_1^v]$, dashed ones represent $E[U_1^v] - \frac{1}{2\psi}$ and $E[U_1^v]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

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 $-(x_1-\overline{\alpha})^2$ 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 ($x_V = \overline{\alpha}$). 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 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

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

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

substance, I assume that $\overline{\alpha} \leq x'$, 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: $\overline{\alpha}(2\gamma_V - 1)$. In the second period it will instead reflect her updated beliefs: $\overline{\alpha}(2\mu_V - 1)$ (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:

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

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 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.

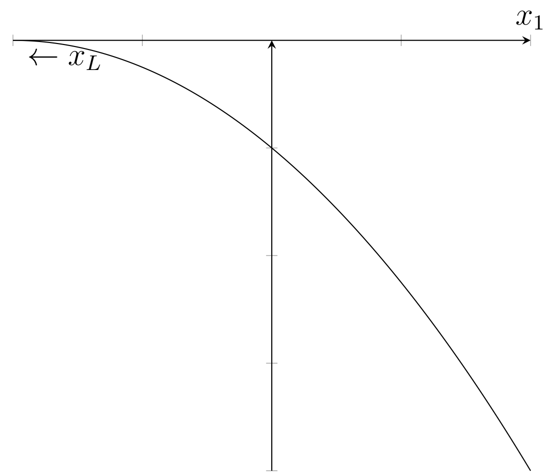

Figure 2: Party L's first-period utility as a function of the implemented policy

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* $\tilde{\tilde{\overline{\alpha}}}$ *and* $\widetilde{x_L}$ *such that if* $\overline{\alpha} > \tilde{\tilde{\overline{\alpha}}}$ *and* $x_L < \widetilde{x_L}$*, 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* $[0, \infty]$*.*

The indirect effect is stronger if information has a large impact on the voter's policy preferences: the larger $\overline{\alpha}$, 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

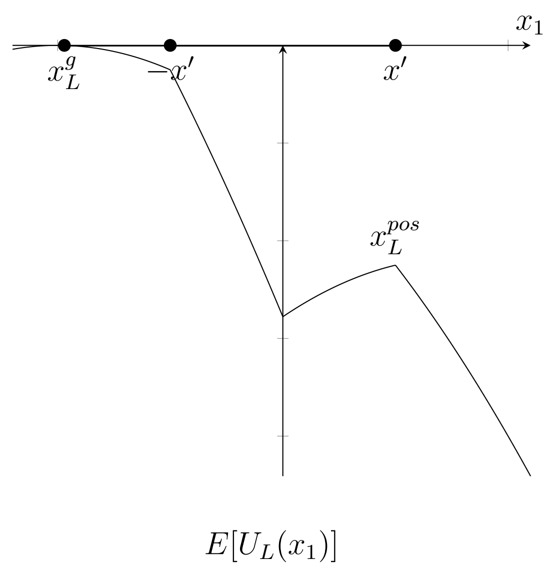

Figure 3: Party L's expected utility as a function of first-period policy

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) $x_R > \overline{\alpha}$, 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 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 ($\overline{\alpha}(2\gamma_V - 1)$). 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

[4]Notice 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).

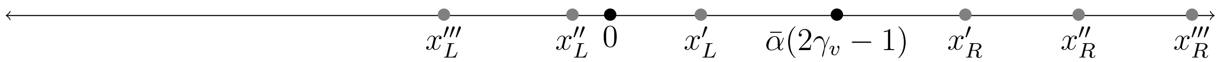

Figure 4: Platforms symmetric around the voter's preferred policy

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

(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 exits. 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* $\widetilde{\widetilde{\alpha}}$ *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:* $\overline{\alpha} > \widetilde{\widetilde{\alpha}}$

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 $x_1 = \overline{\alpha}(2\gamma_V - 1)$ (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 $\overline{\alpha}$ is identical). If the voter receives no additional information, the parties will converge on $\overline{\alpha}(2\gamma_V - 1)$ 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

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

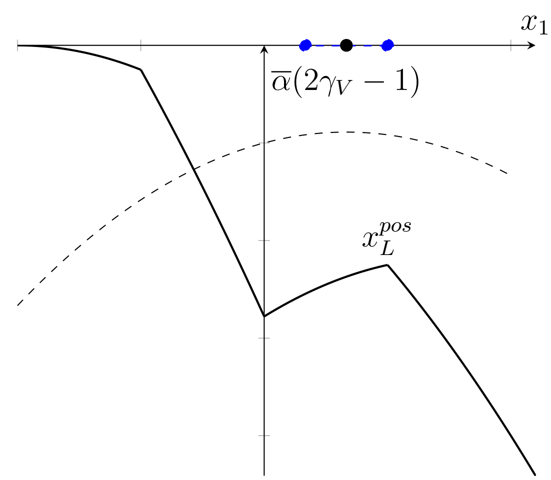

$E[U_L(x_1)], E[U_V(x_1)]$

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.

from a successful gamble is therefore increasing in $\overline{\alpha} = -\underline{\alpha}$. Additionally, the value of moving tomorrow's equilibrium policy increases as the party's bliss point 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.* $\frac{\partial \widehat{x_L}}{\partial \gamma_V} > 0$ *and* $\frac{\partial \widehat{\overline{\alpha}}}{\partial \gamma_V} < 0$*)*

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. $x_1^{R^*} - \overline{\alpha}(2\gamma_V - 1) = \overline{\alpha}(2\gamma_V - 1) - x_1^{L^*}$;

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 $2\overline{\alpha}(2\gamma_V - 1) - x_L^{pos}$.

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

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

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

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: $x_1^{R^*} = \overline{\alpha}(2\gamma_V - 1) - x_L^{Min}$ and $x_1^{L^*} - \overline{\alpha}(2\gamma_V - 1) = \overline{\alpha}(2\gamma_V - 1) - x_1^{R^*}$, 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* $\tilde{\tilde{\alpha}}$*,* $x_L^\dagger$*,* $x_R^\dagger$ *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:* $\overline{\alpha} > \tilde{\tilde{\alpha}}$
- *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 $\overline{\alpha}$. 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 $\overline{\alpha}$).

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 $\overline{\alpha}(2\gamma_V - 1)$).

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

can be sustained in a gambling equilibrium. Further, if $E[U_L(2\overline{\alpha}(2\gamma_V - 1) - \tilde{x})] \leq E[U_L(\tilde{x})]$ 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

[7] Recall 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.

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 wiling 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 $2\overline{\alpha}(2\gamma_V - 1) - \tilde{x}$ and strictly increase its expected utility. If instead $E[U_L(2\overline{\alpha}(2\gamma_V - 1) - \tilde{x})] \leq E[U_L(\tilde{x})]$, the unpopular party can do nothing better than allow its opponent to win (recall that L's expected utility is monotonically decreasing on $[2\overline{\alpha}(2\gamma_V - 1) - \tilde{x}, 0]$ 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* $\widetilde{\overline{\alpha}}$, 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:* $\overline{\alpha} > \widehat{\overline{\alpha}}$
- *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 $\overline{\alpha}$, 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

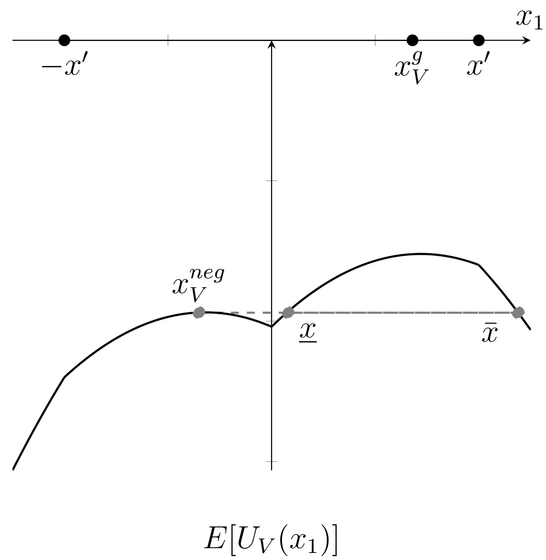

Figure 6: Forward looking voter's expected utility as a function of first-period policy

weakly better off. This makes it hard to identify pairs of platforms such that the left-wing party has no profitable deviation.

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 $\overline{x} > x_V^g$ such that $E[U_V(\underline{x})] = E[U_V(\overline{x})] = E[U_V(x_V^{neg})]$, and $E[U_V(x)] > E[U_V(x_V^{neg})]$ for any $x \in (\underline{x}, \overline{x})$ (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. $\underline{x} \leq x_1^{L^*} \leq x_V^g \leq x_1^{R^*} \leq min \in \{\overline{x}, x_L^{pos}, \tilde{x}\}$, *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', textitThe 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 Jour- nal of Economics*, 119: 2, 223?48.

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,

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

Given the assumption that ϵ is uniformly distributed

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

Therefore the above simplifies to

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

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 $|x_1| \geq |x'|$*

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

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

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

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

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

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

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

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

$$x\geq\frac{1}{4\bar{\alpha}\psi}=x' \tag{10}$$

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

$$Prob(L=1|\gamma,0<x_1<x')=$$

$$\gamma[Prob(-(x_t-\bar{\alpha})^2+e_1>-(x_t-\underline{\alpha})^2+\tfrac{1}{2\psi})]+(1-\gamma)[Prob(-(x_t-\underline{\alpha})^2+e_1<-(x_t-\bar{\alpha})^2-\tfrac{1}{2\psi})] \tag{11}$$

Given the symmetry

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

(15) simplifies to

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

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:

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

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:

$$\begin{aligned} &\underset{x_1}{\text{maximise}} \quad -(\mathrm{x_1} - \mathrm{x_L})^2 - (1 + 4\overline{\alpha}\psi\mathrm{x_1})(\overline{\alpha}(2\gamma_\mathrm{V} - 1) - \mathrm{x_L})^2 + 4\overline{\alpha}\psi\mathrm{x_1}(\underline{\alpha} - \mathrm{x_L})^2 \\ &\text{subject to} \quad x_1 \in [-\frac{1}{4\overline{\alpha}\psi}, 0] \end{aligned} \tag{15}$$

- 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* $\widetilde{\widetilde{\alpha}}$ *and* $\widetilde{x_L}$ *such that if* $\overline{\alpha} > \widetilde{\widetilde{\alpha}}$ *and* $x_L < \widetilde{x_L}$ *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:

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

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

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

and

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

□

Proposition 1: *The exist unique $\widehat{x_L} \in (0, \widetilde{x_L})$ and $\tilde{\tilde{\alpha}}$ 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: $\overline{\alpha} > \tilde{\tilde{\overline{\alpha}}}$*

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:

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

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

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

And

$$\bar{\alpha}^2 > \frac{1}{8\psi\gamma_V} \tag{21}$$

The proof of Corollary 1 follows straightforwardly from above. □

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. $x_1^{R^*} - \overline{\alpha}(2\gamma_V - 1) = \overline{\alpha}(2\gamma_V - 1) - x_1^{L^*}$;

2. $x_1^{L^*} \geq x_L^{Min}$

Proof. The proof for Points 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* $E[U_L(\hat{x})] = E[U_L(2\overline{\alpha}(2\gamma_V - 1) - \hat{x})]$

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)* $E[U_L(\hat{x})] = E[U_L(2\overline{\alpha}(2\gamma_V - 1) - \hat{x})]$, *(ii) for any* $x < \hat{x}$, $E[U_L(x)] > E[U_L(2\overline{\alpha}(2\gamma_V - 1) - x)]$ *and (iii) for any* $\hat{x} < x < 0$, $E[U_L(x)] < E[U_L(2\overline{\alpha}(2\gamma_V - 1) - x)]$.

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

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

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

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

When $x \leq -x'$ (since both x and $2\overline{\alpha}(2\gamma_V - 1) - x$ 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

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

The uniqueness of $\hat{x}$ follows straightforwardly from the fact that $E[U_L(2\overline{\alpha}(2\gamma_V - 1) - x]$ 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 $[2\overline{\alpha}(2\gamma_V - 1) - x_L^{pos}, 0]$[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 $2\overline{\alpha}(2\gamma_V - 1) - x_L^{pos}$. 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 conclues 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:

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

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

[9] When 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.

Which reduces to:

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

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

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

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* $\tilde{\overline{\alpha}}$, $x_L^\dagger$, $x_R^\dagger$ *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* : $\overline{\alpha} > \tilde{\overline{\alpha}}$
- *The parties are sufficiently extreme:* $x_L < x_L^\dagger$ *and* $x_R > x_R^\dagger$
- *The parties are sufficiently ideological in their beliefs*: $\gamma_L < \gamma < \gamma_R$

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

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

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

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

$$\gamma_L < \gamma_V - \frac{1}{8\overline{\alpha}^2\psi} \tag{30}$$

and

$$\overline{\alpha}^2 > \frac{1}{8\psi\gamma_V} \tag{31}$$

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

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

Thus $\frac{\partial E[U_R(x_1)]}{\partial x_1} = -2(x_1 - x_R) + 4\overline{\alpha}\psi(2\overline{\alpha}(2\gamma_V - 1) - x_R)^2 - 4\overline{\alpha}\psi(\gamma_R(\overline{\alpha} - x_R)^2 + (1 - \gamma_R)(\underline{\alpha} - x_R)^2)$. The equilibrium condition is therefore:

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

Which can be rewritten as:

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

Which requires

$$\gamma_R > \gamma_V - \frac{1}{8\overline{\alpha}^2\psi} \tag{34}$$

□

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. $x_1^{R^*} - \overline{\alpha}(2\gamma_V - 1) = \overline{\alpha}(2\gamma_V - 1) - x_1^{L^*}$

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:

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

1. $x_1^{R^*} - \overline{\alpha}(2\gamma_V - 1) = \overline{\alpha}(2\gamma_V - 1) - x_1^{L^*}$

2. $x_1^{L^*} \geq \widetilde{x_L^{Min}}$

can be sustained in a gambling equilibrium. Further, if $E[U_L(2\overline{\alpha}(2\gamma_V - 1) - \tilde{x})] \leq E[U_L(\tilde{x})]$ *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* $\widetilde{\widetilde{\alpha}}$, 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:* $\overline{\alpha} > \widetilde{\widetilde{\alpha}}$)
- *The parties are sufficiently ideological in their beliefs:* $\gamma_L < gamma < \gamma_R$
- *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:

$$\begin{aligned} \underset{x_1}{\text{maximise}} \quad & -\gamma_{\mathrm{V}}(\mathrm{x}_1 - \overline{\alpha})^2 - (1-\gamma_{\mathrm{V}})(\mathrm{x}_1 - \underline{\alpha})^2 - \delta(1 - 4\overline{\alpha}\psi\mathrm{x}_1)[\gamma_{\mathrm{V}}(\overline{\alpha}(2\gamma_{\mathrm{V}} - 1) - \overline{\alpha})^2 + (1-\gamma_{\mathrm{V}})(\overline{\alpha} \\ \text{subject to} \quad & x_1 \leq \frac{1}{4\overline{\alpha}\psi} \end{aligned}$$

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

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

And

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

These reduce to

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

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

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

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

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

□

Proposition 6: *Any pair of platforms satisfying:*

1. $E[U_V(x_1^{L^*})] = E[U_V(x_1^{R^*})]$
2. $\underline{x} \leq x_1^{L^*} \leq x_V^g \leq x_1^{R^*} \leq min \in \{\overline{x}, x_L^{pos}, \tilde{x}\}$, *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. □

Do We Get the Best Candidates When We Need Them the Most?

Abstract

Do the right candidates for office choose to run at the right time? I analyze a model of repeated elections in which politicians differ in the probability of being competent. Voters update their beliefs about the office holder's ability upon observing his performance in office. In each period, the country faces either a safe situation or a crisis. A crisis has two key features: it exacerbates the importance of the office holder's competence and, as a consequence, the informativeness of his performance. I show that electoral accountability has the perverse consequence of discouraging good candidates from running in times of crisis. Precisely when the voter would need him the most, the politician who is most likely to be competent chooses to stay out of the race in order preserve his electoral capital. In contrast with results in the existing literature, this adverse selection emerges even if running is costless and if office is more valuable than the outside option.

Introduction

A growing empirical literature highlights that the quality of political leaders has a critical impact on a country's performance (e.g. Jones and Olken 2005, Besley, Montalvo and Reynal-Querol, 2011). From a theoretical perspective, it then becomes essential to understand under which conditions high-quality politicians are willing to run for office in the first place. One question is particularly important to evaluate the effectiveness of democratic elections in improving voters' welfare: do the right candidates self-select at the right time? More specifically, are the most competent politicians willing to run for office during times of crisis, when competence matters the most?

The formal literature has so far placed little emphasis on this question. Most extant models of elections in fact take the pool of candidates as exogenous, focusing instead on voters' ability to identify good politicians to be (re)elected and bad ones to be thrown out. A small recent literature allows for endogenous candidate entry, thereby analysing the equilibrium supply of good politicians. However, these works typically consider a static setting, focusing on *which types* self-select into politics and highlighting the difficulty of attracting competent politicians if office rents are too low compared to private market salaries. Little attention is instead paid to *when* the right candidates are willing to run, if a longer planning horizon is considered.

In this paper, I adopt a very different perspective. I consider a world in which potential candidates are career politicians, for whom office is always more valuable than the outside option. As such, entering the race is always the statically optimal choice for all potential candidates, irrespective of their expected ability and the conditions in the country. I show that this does not always hold true when we take into account their dynamic incentives. Under some conditions, 'good' candidates are not willing to run for office during times of crisis. The politician who is most likely to solve the crisis also has the most to lose from failing. As such, precisely when the voter would need him the most, he chooses to stay out of the race in order to preserve his electoral capital for the future. In contrast, the 'worst' (in expectation) potential candidate is always willing to take the gamble, and run for office during challenging times. Thus, the voters gets the *wrong candidates* at the *wrong time*. Crucially, this adverse selection does not arise due to weak electoral incentives, as it is the case the extant literature. Quite the opposite, it emerges precisely as a perverse consequence of accountability.

In the baseline model I consider a game with two time periods and an election in each. The players are two potential candidates and a representative voter. Potential candidates are career politicians, that differ from each other in the probability of being competent (a politician's true type is unknown to all players). In each period, the politicians simultaneously choose whether to run for office. The model is one of pure selection: the office holder's performance results in either a good or a bad governance outcome, with the probability of producing a good outcome a function of the incumbent's true type. Politicians are office motivated, an their (per-period) payoff from holding office is always higher than their outside option. This payoff consists of both monetary and ego-rents: while monetary rents are always accrued in the same measure, ego rents represent the legacy payoff that an office holder only enjoys when he delivers a good performance.

The crucial element of the model is that in each period the country either experiences a crisis, or a period of 'business as usual'. A crisis (economic or otherwise) is an *exogenous* shock that has two key features: it amplifies the impact of the office holder's competence and, at the same time, the informativeness of his performance. In other words, precisely because competence matters the most during times of crisis, this is also when the governance outcomes reveals most information about the office holder's ability. In particular, I assume that both competent and incompetent politicians are always able to deliver a good performance during normal times. In contrast, the office holder's ability determines the probability that he would successfully manage a crisis. A crisis therefore provides the voters with a 'test' of the incumbent's quality.

Consider the incentives a career politician faces. In the last period election, a politician must only evaluate the expected value of holding office today. This is always higher than the payoff from staying home, therefore all potential candidates are always willing to enter the race. Not so much in the first period. When politicians choose whether to run for office, they must consider both the expected payoff from being elected today, and how it influences the chances of being elected tomorrow. Suppose that the country is hit by a crisis in the first period. This has two consequences. First, the office holder's performance will reveal information about his true ability, and therefore influence his future electoral prospects. Second, the value of holding office today is lower than the expected rents from being in office the second period: a crisis may not arise again tomorrow, therefore the probability that the office holder will be able to deliver a good performance and enjoy the associated legacy payoff is higher in the second term. In this sense, the first period office holder is taking a gamble. The lower the probability of being competent, the riskier this gamble.

Given the reasoning above, it may seem counter-intuitive that precisely the politician who is most likely to be competent would decide to stay out of the race during challenging times. However, while this politician has the highest chances of surviving a crisis, he also has a valuable electoral advantage. As a consequence, information can only hurt his future electoral chances, and this politician experiences *fear of failure*. Under some conditions, he will therefore choose to stay out of the race when a crisis is likely to arise, so as to prevent the voter from learning about his true ability. In contrast, the worst (in expectation) potential candidate never has anything to lose from holding office in the first period. Indeed, holding office during times of crisis can only increase his future electoral chances, by allowing him to prove himself. As such, he always has incentives to *gamble on his own success*, and is willing to enter the race under both states of the world. Thus (under some conditions) only the worst candidate is willing to run for office during challenging times.

The second result of the paper shows that, by influencing the pool of candidates that self-select in equilibrium, an exogenous shock can also impact the electoral effect of incumbency. I consider a setting in which incumbency does not provide any advantage (or disadvantage) in terms of campaign resources or name recognition. Nonetheless, while incumbency status has no effect on electoral performance during periods of 'business as usual', an office holder that runs for re-election in times of crisis may experience either an advantage or a disadvantage.

In a robustness section I analyse several variants of the model, relaxing some of the most restrictive assumptions imposed in the baseline set-up (in a separate section I also introduce an amended version of the game, in which politicians live for more than two periods but only care about the monetary rents from office). I show that while the adverse selection documented in this paper can me more or less severe, it is unlikely that any democracy may be immune from it. Indeed, the source of this inefficiency seems to lie precisely in the accountability relationship between the voters and their representatives. The problem that the voters face is that they cannot credibly commit to ignoring valuable information that may be generated about the incumbent. Precisely when competence matters the most, the office holder's performance reveals most information about his true ability. The politician who is most likely to be competent also has the most to lose, and is unwilling to take the gamble. As a consequence, the more the voters would want a competent politician in office, the less likely they are to get one. This highlights how the nature of the results presented in this paper differs from the findings in the extant literature, where adverse selection instead emerges due to low powered electoral incentives.

Literature Review

This paper contributes to the literature on the endogenous supply of good politicians (Caselli and Morelli 2004, Messner and Polborn 2004, Besley 2005, Dal Bo, Dal Bo and Di Tella 2006, Mattozzi and Merlo 2008, Fedele and Natticchioni 2013, Brollo 2013).[1] This literature has so far focused mainly on how an individual's outside option in the private market influences his decision to run for office. Political ability and private market salary are assumed to be correlated, therefore good politicians also have a higher opportunity cost of holding office. This potentially generates an adverse selection, whereby low ability individuals are more likely to enter politics.

[1] Other scholars analyse endogenous entry, but focus on settings in which potential candidates differ in motivations (see Callander 2008) or ideology (see Osborne and Slivinski 1996, Besley and Coate, 1997), rather than quality.

As highlighted above, this paper adopts a completely different perspective. It considers a world in which potential candidates are career politicians, for whom the value of holding office is always higher than the expected payoff from the private market. Thus, rather than looking at the financial considerations that drive self-selection into politics, I focus on how politicians' dynamic electoral incentives influence the *timing* of their entry decision.

The crucial feature of this model is to allow exogenous shocks to the country's conditions to influence the *endogenous* opportunity cost of holding office. As such, this paper is in close conversation with a recent literature in formal theory, that highlights how events outside of the office holders' control may nonetheless impact their electoral fortunes, by altering the inferences voters draw upon observing their performance in office (see Ashworth, Bueno de Mesquita and Friedenberg, 2017 and 2018). These works complement the model presented here, since they take the pool of candidates as given and focus instead on how crises influence office holders' effort choice.

Finally, this model connects with several papers that analyse political actors' incentives to gamble, within the framework of a multi-armed bandit model (e.g. Strulovici 2009, Dewan and Hortala-Vallve 2018). In these works, political actors must choose between a risky and a safe policy. The consequences of a risky choice inform voters and politicians about the underlying state of the world, or the office holder's true ability. In contrast, the outcome of a safe policy reveals no additional information. The crucial assumption is therefore that office holders are always free to *choose* to generate more or less information. In this paper, I instead assume that the informativeness of governance outcomes is determined exogenously by the 'riskiness' of the situation the country faces. Politicians cannot choose which arm of the bandit to pull, they can only choose whether to play.

The Baseline Model

I study the endogenous supply of competent candidates by analysing a game of repeated elections with two time periods. At the beginning of the game, each party $P \in \{1, 2\}$

draws one potential candidate C^P from the pool of its members. Politicians differ in the probability of being competent. Specifically, each politician is one of two types, good or bad: $\theta_i \in \{G, B\} \; \forall i \in \{C^1, C^2\}$. A politician's type is unknown to all players, including the politician himself. This reflects the assumption that political ability is more than the product of a pre-determined and identifiable skill-set. As such, it can never be verified ex-ante, but only discovered via experience. Players share common beliefs that politician C^P is a good type with probability q^P (formally, party P draws from a pool containing a proportion q_P of good types). I will assume that $q_1 > q_2$. I will therefore refer to C^1 as the ex-ante advantaged potential candidate, and to C^2 as the disadvantaged one.

At the beginning of each period, the two potential candidates simultaneously choose whether to run for office. If C^P chooses to stay out of the race, party P is unable to field a viable candidate and it resorts to a reserve candidate R^P, which is known to be a bad type with probability one (this assumption is without loss of generality). The existence of the reserve candidates R^1 and R^2 is imposed for purposes of presentation in order to avoid equilibria with uncontested elections, but otherwise has no effect on the results. Once the candidates are endogenously determined, a representative voter V chooses whom to elect.

In each period, the country either faces a normal situation, or it is hit by a negative shock: $\omega_t \in \{N, S\} \; \forall t \in \{1, 2\}$. A shock is an *exogenous* crisis: it may represent a period of economic hardship, a war, or even a natural disaster. The key feature of a shock is that it amplifies the effect of the office holder's type on his performance: competence matters the most during times of crisis. Specifically, in each period the office holder produces either a good or a bad governance outcome $o_t \in \{g, b\}$, $\forall t \in \{1, 2\}$. The governance outcome is a good one whenever a crisis does not arise, or if it arises but the office holder is able to solve it. Otherwise, the outcome is a bad one. The office holder's type determines the probability that he is able to solve a crisis. A good type always produces a good outcome under a negative shock, whereas a bad one does so with probability $\beta \in [0, 1]$:

- $prob(o_t = g|\omega_t = N, \theta_t = G) = prob(o_t = g|\omega_t = N, \theta_t = B) = 1$

- $prob(o_t = g|\omega_t = S, \theta_t = G) = 1$
- $prob(o_t = g|\omega_t = S, \theta_t = B) = \beta < 1$

This specific parametrization is adopted for simplicity, but is not necessary for the results. Notice that the parameter β can be interpreted as the complexity of the crisis, but also as the country's resiliency. For example, when a country can count on a competent bureaucratic apparatus, it is more likely to survive a negative shock even if an incompetent type is in office.

Arguably, there are substantive reasons to defend the assumption that competence matters the most in times of crisis. However, it is also important to highlight that if we allow players to live for more than two periods the key qualitative result of the paper (i.e., the voter is less likely to get the best candidate precisely when she needs him the most) would continue to hold under the opposite assumption, that is if crises mute rather than amplifying the impact of the office holder's type. I will discuss this further in a separate section.

Without loss of generality, I assume that the state of the world ω_t realizes after the election has taken place. Players share common prior beliefs that $prob(\omega_t = S) = \bar{p}$, with ω_t i.i.d. in each period. In a robustness section, I relax this assumption and allow the probability of a crisis in the second period to be a function of the first period incumbent's performance. At the beginning of each period, players also observe a public signal indicating the likelihood of a crisis arising during the upcoming term. Formally, players observe a signal $\chi \in \{N, S\}$, accurate with probability $\psi > \frac{1}{2}$ ($prob(\chi_t = S|\omega_t = S) = prob(\chi_t = N|\omega_t = N) = \psi > \frac{1}{2}$).

Finally, we must specify the players' payoffs. The voter cares about governance outcomes. She pays a cost $\lambda > 0$ in each period in which $o_t = b$, whereas the payoff from a good outcome $o_t = g$ is normalized to 0. Politicians are office motivated. The value of holding office has two components: monetary rents $K > 0$ and legacy payoffs $\gamma > 0$. While the monetary rents are always accrued by the office holder, the legacy payoffs are conditional on delivering a good performance.[2] γ may represent the 'warm glow feeling' politicians experience when they produce a good governance outcome, or (in a reduced-form) the instrumental value of a good performance. Since the aim of this paper is to focus on the *endogenous* opportunity cost of office, I assume that a politician's outside option is always lower than his per-period payoff from being in office. Politicians' utility when out of office is therefore normalized to 0. Finally, since the focus of this paper is on politicians' incentives and disincentives to *hold* office, I assume that running is costless. However, because I consider a deterministic election process, this assumption has no impact on the qualitative results other than avoiding equilibria with uncontested elections.

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

1. Nature draws the potential candidates' types $\theta_{C^1}, \theta_{C^2} \in \{G, B\}$ and the first period state of the world $\omega_1 \in \{N, S\}$
2. The players observe a public signal $\chi_1 \in \{N, S\}$, accurate with probability ψ
3. C^1 and C^2 simultaneously choose whether to run. If party $P \in \{1, 2\}$ is unable to field a viable candidate it resorts to the reserve R^P.
4. The voter decides whom to elect
5. The first period state of the word ω_1 realizes
6. The fist period governance outcome o_1 realizes
7. The second period starts and nature draws $\omega_2 \in \{N, S\}$
8. The game proceeds as above

To avoid trivialities, I will exclude equilibria in weakly dominated strategies. Since running for office is costless, this implies that a politicians' entry decision is conditional on winning the election.

[2]In a two-period setting the assumption that $\gamma > 0$ is necessary to obtain the results. In a separate section I consider a longer time horizon, and I show that the inefficiency documented in the baseline model survives even if the office holder's payoff is not a function of his performance.

Notice that this is a model of pure selection: the office holder's performance is determined by his true ability and the state of the world, and I do not allow politicians to invest in (costly) effort in order to improve their expected performance and electoral chances. The choice to abstract from this moral hazard problem is purely for presentation purposes and, as long as the governance outcome remains informative at all levels of effort, relaxing this assumption would not alter the main message of the paper.

Analysis

In this section I solve the two-period game and identify the conditions under which adverse selection emerges in equilibrium. Consider first the voter's electoral decision. The voter cares exclusively about governance outcomes. In each period, she therefore elects the candidate that is most likely to deliver a good performance. Straightforwardly, her first period electoral choice is simply a function of her prior beliefs over the candidates' ability. In contrast, the incumbent's performance informs the voter choice in the second period election. This paper builds on a key intuition: the inferences that voters draw upon observing the governance outcome are a function of the state of the world. Thus, the same outcome may convey different information under different environment conditions. In other words, crises have an informational value. Precisely because crises amplify the effect of competence on outcomes (i.e., for any outcome $o_t \in \{g, b\}$, $|prob(o_t|\omega_t = S, \theta_t = G) - prob(o_t|\omega_t = S, \theta_t = B)| > |prob(o_t|\omega_t = N, \theta_t = G) - prob(o_t|\omega_t = N, \theta_t = B)|$), they also increase the informativeness of the incumbent's performance: when the country is hit by a negative shock, the voter is able to draw more precise inferences on the office holder's type. In particular, given the specific parametrization adopted here, both types are always able to deliver a good outcome under a normal state of the world, therefore the office holder's performance is completely uninformative. In contrast, an exogenous crisis provides the voter with a 'test' of the incumbent's political ability, and therefore an opportunity to learn. Denote as $\mu_i(I_i, \omega_1, o_1)$ the posterior probability that politician i is a good type, given his incumbency status $I_i \in \{I, \emptyset\}$, the first period outcome and state of the world. Recall that q_i is the prior probability that politician i is a good type. The following Lemma holds:

Lemma 1. *Suppose that $\omega_1 = N$. Then, the incumbent's performance reveals no information about his type, and the voter's posterior is always equal to her prior beliefs. Suppose instead that $\omega_1 = S$. Then, the voter always obtains new information: for all outcomes $o_1 \in \{g, b\}$ and all politicians $i \in \{C^1, C^2\}$, $\mu_i(I, S, o_1) \neq q_i$.*

This implies that even if a shock is fully exogenous, it may influence the incumbent's electoral chances. Indeed, the voter's decision in the second period may be different under different states of the world, even fixing the governance outcome. In particular, both C^1 and C^2 would be ousted after producing a bad outcome and would be re-elected after producing a good outcome under a crisis.[3] However, a good performance during normal times always guarantees C^1's survival, but is never enough for the ex-ante disadvantaged C^2 to get re-elected.

With this in mind, let us now focus on the potential candidates' incentives. As highlighted above, the model considers a world in which potential candidates are career politicians, for whom the expected *per-period* value of holding office is always higher than the outside option $(K + \gamma[1 - prob(\omega_t = S|\chi_t) + prob(\omega_t = S|\chi_t)(q_i + (1 - q_i)\beta)] \geq K > 0)$. Further, recall that I assume running to be costless. Absent any future electoral considerations, it is therefore straightforward to verify that both viable candidates C^1 and C^2 always have a dominant strategy to run for office in the second period. Excluding equilibria in weakly dominated strategies, the following holds:

Lemma 2. *In any equilibrium of the game, both viable candidates C^1 and C^2 choose to run for office in the second period.*

Not so much in the first period. When choosing whether to run or stay out of the race, politicians consider both the expected value of holding office today and, given Lemma 1, how it influences the probability of being elected tomorrow (i.e., the *endogenous* opportunity cost of office). Both are a function of the state of the world. The

[3] To avoid trivialities, I assume that $\mu_{C^2}(I, S, g) > q_1$, where $\mu_{C^2}(I_{C^2}, \omega_1, o_1)$ is the posterior probability that C^2 is a good type, given his incumbency status $I_{C^2} \in \{I, \emptyset\}$, the first period outcome and state of the world.

per-period expected value of office is always lower in times of crisis ($\omega_1 = S$), since a politician who turns out to be incompetent may be unable to deliver a good outcome and enjoy the associated legacy payoffs. Consider instead the opportunity cost of holding office in the first period. Under a normal state of the world ($\omega_1 = N$) the voter will obtain no new information upon observing the governance outcome (the voter's posterior on the incumbent's type is always equal to her prior). Therefore, holding office today does not influence the probability of being elected tomorrow. In contrast, if a crisis arises the incumbent's performance will reveal information about his true ability. The office holder then risks exposing himself as a bad type and losing the second period election.

Given the above reasoning, it follows straightforwardly that politicians have no reason to stay out of the race when $\chi_1 = N$. The public signal indicates that a crisis is unlikely to arise during the first term, more precisely that a crisis today is less likely than a crisis tomorrow (recall that, given the martingale property of beliefs, the expected posterior probability of a shock in the second period is always equal to the prior $\bar{p}$). As such, the expected rents from holding office today are higher than the expected value of office in the future. Then, irrespective of the opportunity cost in terms of future electoral chances, both potential candidates always choose to enter the race when $\chi_1 = N$. Suppose instead that the public signal indicates that a crisis is likely to arise $\chi_1 = S$. Now, holding office in the future is, in expectation, more valuable. A potential candidate may therefore be worried that, if the crisis materializes, his performance in office would expose him as an incompetent type and hurt his electoral chances in the second period. Straightforwardly, this risk is higher the lower the probability of being a good type. This reasoning may lead us to conclude that positive selection emerges in equilibrium: the politician is most likely to be able to solve a crisis has the strongest incentives to run. Instead, the analysis shows that the opposite is true:

Proposition 1. *In equilibrium, both C^1 and C^2 always run for office under $\chi_1 = N$. Consider instead $\chi_1 = S$. Suppose that the following conditions are satisfied:*

(i) The public signal is sufficiently accurate ($\psi > \underline{\psi}$)

(ii) C^2 is sufficiently unlikely to be a good type ($q_2 < \bar{q}_2$)

(iii) A bad type is sufficiently unlikely to deliver a good outcome under a crisis ($\beta < \bar{\beta}$)

Then, there exists an interval $[q_2, \bar{q}]$ such that when $q_1 \in [q_2, \bar{q}]$, C^1 chooses to stay out and Party 1 resorts to the reserve candidate R^1. Instead, C^2 always chooses to enter the race.

Proposition 1 presents a very stark inefficiency result: in equilibrium, the voter gets the wrong candidates at the wrong time. The ex-ante disadvantaged C^2, which has the lowest expected quality, is always willing to run for office. Instead, it is the politician who is most likely to be competent that sometimes chooses to stay out of the race. To make matters even worse, he does so precisely when the voter would need him the most: the country is very likely to experience a crisis (the public signal is negative and sufficiently informative), competence really matters in times of crisis ($\beta < \bar{\beta}$), and the alternative is really bad ($q_2 < \bar{q_2}$).

To understand this result, let us focus first on the strategic incentives faced by the disadvantaged C^2. Straightforwardly, C^2 would always lose the first period election if C^1 chooses to enter the race. Since running is costless, C^2 is indifferent between entering the race and staying out. Suppose instead that C^1 chooses to sit the first period election out. Now, C^2 must consider how holding office today would influence the probability of being elected tomorrow. Perhaps counter intuitively, holding office during times of crisis would always improve C^2's future electoral prospects, irrespective of how unlikely he is to be able to deliver a good governance outcome. C^2 will only win the second period election if the voter updates positively about his type, or negatively about C^1's ability. If C^1 stays out of the race in the first period, the voter will obtain no new information about his competence. As such, C^2 will always lose tomorrow's election if he chooses to stay home today. The only way to improve his future electoral prospects is to prove himself: prove able to deliver a good governance outcome even after being hit by a negative shock. In other words, the ex-ante disadvantaged politician never has anything to lose from holding office in times of crisis, since new information can only increase his future expected payoff. Running for office in the first period therefore always weakly increases both his immediate and future expected payoff. Thus, irrespective of how likely a crisis is to arise, and how unlikely he is to be able to solve it, C^2 always has incentives to *gamble on his own success*, and has a weakly dominant strategy to enter the race under both realizations of the public signal.[4]

The ex-ante advantaged C^1 faces very different incentives. He is more likely to be able to solve a crisis if it arises, and deliver a good governance outcome. He therefore has a higher expected payoff from holding office today, and a higher likelihood of being re-elected tomorrow. However, C^1 also has a valuable electoral advantage that he does not want to waste. Because of this advantage, information can only hurt his future electoral performance: if the voter learns nothing new, C^1 always wins for sure in the second period. As a consequence, he would want to prevent the voter from learning about his true ability so as to maximize his future electoral chances. In other words, C^1 experiences *fear of failure*: he has incentives to avoid a gamble, even if it is likely to succeed. Therefore, when the public signal indicates that a crisis is likely to arise in the first period, C^1 faces a trade-off. If he chooses to stay out of the race, his immediate payoff decreases as he foregoes the rents from holding office today. However, if he chooses to run, he risks exposing himself as a low type and therefore wasting his electoral capital and losing tomorrow, when holding office is in expectation more valuable. The problem that he faces is that there is no safe strategy. If he chooses to run, he gambles on his own success. That is, on the probability of being able to deliver a good performance even under a crisis. If he chooses not to run, he gambles on his opponents failure. That is, on the probability that if a crisis arises C^2 will not be able to solve it and win re-election in the second period. C^1's equilibrium choice will therefore depend on the expected value of holding office today versus tomorrow, and on the relative riskiness of the two gambles. The equilibrium conditions are intuitive. When a crisis is very likely, C^2 is unlikely to reveal himself as a good type, and C^1 is not sufficiently confident in his own ability, he chooses to stay out of the race so as to preserve his electoral capital for the future.

[4]Notice that the same holds for the reserve candidates R^1 and R^2, who are therefore always willing to represent their respective party in the general election.

In concluding this section it is important to emphasize that the nature of the inefficiency documented in Proposition 1 is very different from analogous results presented in the literature. Extant works highlight the difficulty of attracting good politicians if office rents are too low to compensate for their outside option in the private market. In other words, adverse selection emerges due to weak electoral incentives. Here, the opposite is true. In this model, running is costless and holding office is always more valuable than the outside option. The inefficiency documented above emerges precisely as a perverse consequence of electoral accountability. The problem that the voter faces is that she can never credibly commit to ignoring valuable information that may be revealed about the incumbent. Precisely because competence matters the most in times of crisis, this is also when governance outcomes are most informative. The politician who is most likely to survive a crisis is also the one who has the most to lose, and is therefore unwilling to take the risk. As such, these results speak to an open debate in the literature: is voter competence is actually good for voters? Scholars have argued that a rational and more informed electorate may paradoxically induce office holders to exert less effort, or adopt worse policies (see Ashworth et al. 2014). This paper suggests that the problem may run even deeper: voters' inability to commit to ignoring information about the incumbent's performance may prevent them from attracting competent politicians to office in the first place.

Discussion and Robustness

For the purpose of simplifying the presentation and thus focusing on the key intuition underlying the results, the model analysed here considers a stylized environment with a binary state of the world and governance outcome, and imposes parameter values such that outcomes are only informative during periods of crisis. These are very stark assumptions, but are not necessary for the emergence of the inefficiency documented above. As highlighted by the discussion in the previous section, the key property of the model that underpins the results is that crises amplify the impact of the office holder's type *and*, at the same time, the informativeness of his performance. Governance outcomes are most informative precisely when competence matters the most.

Ashworth et al. (2017) show that this property holds more generally, even under a less stylized information environment. They look at a world in which, similarly to the model presented here, governance outcomes are the output of a production function that depends on the incumbent's type and two shocks: the observable disaster (i.e., the state of the world) and an unobservable idiosyncratic shock. Their general model then backs away from any specific functional form assumption. Their results thus hold for any production function increasing in type and decreasing in disaster intensity, any distribution of the state of the world, and any (symmetric) distribution of the random shock satisfying a strict monotonic likelihood ratio property. For our purposes, their key finding is to show that 'governance outcomes are more informative (resp. less informative) following larger disasters, if disasters amplify (resp. mute) the effect of type' (2017, p. 12). In other words, exactly as in the stylized setting considered here, outcomes are most informative when competence matters the most. It is crucial to stress that *Ashworth et al. (2017) do not allow for endogenous candidate entry* (indeed, in their model politicians are dummies, that do not take any strategic action). As such, their work should be considered as a complement to this paper, indicating that the inefficiency highlighted in Proposition 1 holds beyond the specific assumptions adopted here.

A second assumption imposed in the model is that governance outcomes influence a politician's payoff only when in office. Intuitively, relaxing this assumption will mitigate the adverse selection documented above. However, as I show below, the inefficiency is never eliminated altogether. In the following paragraphs I introduce several variants of the baseline model and informally discuss the results' robustness. All the formal proofs are in Appendix B.

There are several ways in which the office holder's poor performance may negatively affect the other potential candidates' payoffs. First, we may argue that governance outcomes *directly* influence politicians' utility even when they are out of office. Suppose then that politicians, just like the voter, suffer a cost $-\lambda$ whenever a bad governance outcome is produced. Denote $\mathbb{I}_g$ a binary indicator taking value 1 when $o_t = g$, and 0 otherwise. A politician's per period payoff is then $R + \mathbb{I}_g\gamma - (1 - \mathbb{I}_g)\lambda$ when in office, and $-(1-\mathbb{I}_g)\lambda$ otherwise. As in the baseline model, all politicians are always willing to run under $\chi_1 = N$. Similarly, C^2 has no reason to stay out of the race in times of crisis, since holding office always can only increase both his expected payoff today and his electoral chances tomorrow. Consider now the problem that the ex-ante advantaged C^1 faces. Straightforwardly, his incentives to run are higher than in the baseline model. If he chooses to stay out of the race, and free-ride on his opponent, he increases the risk of incurring the cost of of a poor governance outcome. We may be tempted to conclude that, for a sufficiently large λ, C^1 would always be willing to run when C^2 is very likely to be a bad type. Instead, as in the baseline model, the opposite is true. The qualitative results are in fact exactly as indicated in Proposition 1: C^1 chooses to stay out of the race precisely when his opponent is very likely to fail (i.e., q_2 and β sufficiently low). C^1 is willing to increase the risk of suffering the cost λ today, in order to maximise the probability of getting to office tomorrow, when a good performance is easier to deliver. Crucially, this holds for any value of λ. The other comparative statics go in the expected direction: as λ increases, C^1 is more likely to enter the race (in the sense of set inclusion).

Alternatively, just like in the baseline model, we may argue that politicians only care about their own performance in office. Nonetheless, governance outcomes may *indirectly* influence a politician's expected payoff, irrespective of his incumbency status. For example, a bad outcome in the first period may increase the probability of a crisis arising (again) in the second. To account for this possibility, assume that $prob(\omega_2 = S|o_1 = g) = \bar{p}$ and $prob(\omega_2 = S|o_1 = b) = \alpha\bar{p}$, where $\alpha \in (1, \frac{1}{\bar{p}})$. As above, free-riding now comes with a cost for C^1: a bad outcome today decreases the expected value of holding office tomorrow. This tends to increase C^1's incentives to run, but does not alter the conclusions from the baseline model: C^1 chooses to stay out of the race precisely when his opponent is most likely to deliver a poor performance. Importantly, this holds even if a bad outcome in the first period pushes the probability of a future crisis arbitrarily close to one (i.e., α is arbitrarily close to $\frac{1}{\bar{p}}$). A similar reasoning applies if we assume that crises are always exogenous (i.e., the probability that $\omega_2 = S$ is not a function of o_1), but a bad governance outcome decreases the country's future resiliency (β). In other words, the first-period office holder's poor performance reduces the probability that the country would survive a future shock if an incompetent type is in power. [5]

Finally, the baseline model assumes that the office holder always obtains the same payoff from a good performance, irrespective of the state of the world. However, we could argue that producing a good governance outcome under a crisis should yield a higher legacy payoff than performing well during normal times. Suppose then that the office holder's legacy payoff is $\nu(\omega_t)\gamma$, where $\nu(N) = 1$ and $\nu(S) > 1$. Straightforwardly, for a sufficiently large $\nu(S)$, C^1's expected overall payoff from entering the race in the first period is increasing in the probability of a crisis. Perhaps more surprisingly, the likelihood that he chooses to run (in the sense of set inclusion) never is. Recall that C^1 is always guaranteed re-election if he gets to office during normal times. Irrespective of how large is the legacy payoff from solving a crisis, increasing the probability of a shock can therefore only reduce the likelihood that C^1 stands for office in the first period. Thus, the assumption that office holders would obtain a larger legacy payoff in times of crisis alleviates the inefficiency documented above, but does not alter the quality of the results: the more the voter needs a competent politician in office, the less likely she is to get one.

This section has highlighted that the crucial inefficiency identified in Proposition 1 can be more or less severe, but it is unlikely that any democracy may be immune from it. Indeed, this inefficiency seems to lie at the very core of the accountability relationship between voters and politicians.

The Electoral Effect of Incumbency

The results in Proposition 1 indicate that exogenous crises influence the pool of candidates that are willing to run in equilibrium. In the baseline model, I consider an open seat election. However, if we think about an incumbent running for another term, a

[5] Formally, $prob(o_2 = g|\omega_2 = S, \theta_I = B, o_1 = g) = \beta$ and $prob(o_2 = g|\omega_2 = S, \theta_I = B, o_1 = b) = \delta\beta$, where $\delta \in [0, 1]$. θ_I denotes the type of the second period office holder.

question emerges naturally: do exogenous shocks influence the incumbent's electoral chances? In particular, is the electoral effect of incumbency different under different states of the world? In this model, incumbents do not enjoy an exogenous advantage (or disadvantage) in terms of resources or name recognition. In what follows I will also fix the priors on the candidates' ability, so that there is no impact of incumbency status on voters' perception of political competence. I therefore focus exclusively on whether endogenous candidate entry generates an electoral effect of incumbency, and how this changes from times of crisis to periods of business as usual.

To analyse this question, suppose that C^2 is the incumbent office holder at the beginning of the game (so that q_2 is the posterior probability that he is a good type, given the prior and his performance at t=0). Further, suppose that office holders face a term limit of two. Therefore, if C^2 is re-elected in the first period, he cannot run again in the second.[6] The replacement (potential) candidate for Party 2 is then drawn in the second period from a pool with a proportion q_r of good types.

To understand the electoral impact of incumbency, I compare the probability that C^2 wins the first period election in the baseline model (i.e., when the election is open seat) with his first period electoral performance under incumbency status. This is essentially equivalent to comparing C^1's incentives to run in the first period in the two cases. In order to generate continuous probabilities, I assume that q_1 is drawn at the beginning of the game from a uniform distribution on $[q_2, \mu_{C^2}(I, C, g)]$ (recall that I assume $q_1 < \mu_{C^2}(I, S, g)$).

The results show that no effect of incumbency emerges when the players observe a public signal indicating normal times. In contrast, depending on the expected quality of Party 2's replacement candidate, either an incumbency advantage or a disadvantage arises when $\chi_1 = S$. Additionally, irrespective of whether the effect of incumbency is positive or negative, it is always increasing in the signal's accuracy:

Proposition 2. *Incumbency status has no effect on C^2's electoral chances under $\chi_1 = N$. Suppose instead that $\chi_1 = S$. Then, C^2 experiences an incumbency disadvantage*

[6] If no term limits are imposed, the politicians' incentives are exactly as in the baseline model, and incumbency status never has any effect on electoral performance.

whenever $q_r > q_1$*, and an advantage whenever* $q_r < q_1$*. In both cases, the effect of incumbency is increasing in the signal's accuracy* ψ*.*

The first result is straightforward. Irrespective of whether the election is an open seat one, C^1 is always willing to run for office under $\chi_1 = N$. Therefore, C^2 always loses the first period election with probability 1, and incumbency status has no effect on his electoral performance. Suppose instead that a negative signal $\chi_1 = S$ is observed at the beginning of the first period, indicating a crisis is likely to arise. First, let $q_r > q_1$. In this case, C^1 has no electoral capital to preserve for future elections. Indeed, in order to win the second period election he needs the voter to update positively about his type. Thus, C^1 has no reason to stay out of the race, and will always choose to run in equilibrium. This, in turn, generates an incumbency disadvantage: C^2 wins with strictly positive probability in the open seat election, but loses for sure when he runs as the incumbent office holder. This disadvantage increases in the signal's accuracy, since C^1's incentives to run in the open seat election (conditional on $\chi_1 = S$) are weaker the higher the probability of a crisis arising.

Suppose instead that $q_r < q_1$: C^1 always wins the second period election if the voter receives no new information about his type. Here, incumbency status has a positive effect on C^2's electoral performance. To understand this result, consider the incentives C^1 faces in the open seat election. When he chooses not to run for office, C^1 gambles on his opponent's failure. Thus, he must take into account the risk that a crisis arises, and C^2 is actually able to solve it. Conversely, when C^1 must decide whether to run against a term limited incumbent, he does not need to worry about the office holder's expected performance. Indeed, if C^1 stays out of the race today, he always wins tomorrow's election. C^1's incentives to run are stronger in the open seat election, and C^2 experiences an incumbency advantage. Notice that the source of this incumbency advantage is exactly the reverse of the 'scare off' effect typically discussed in the literature (Cox ane Katz 1996, Levitt and Wolfram 1997). C^1 is more likely (in the sense of set inclusion) to stay out of the race precisely because he has nothing to fear from the (term limited) incumbent.

An analogous reasoning explains why this incumbency advantage is increasing in the signal's accuracy ψ. As ψ increases, so does the posterior probability that a crisis will occur in the first period. As a crisis becomes more likely, both C^1's expected payoff from holding office today and his probability of being re-elected tomorrow decrease. Thus, a increase in ψ always has a direct negative effect on C^1's incentives to run. However, in the open seat election an indirect effect also emerges. Recall that C^2 would be re-elected only upon producing a good governance outcome under a crisis. Thus, as ψ increases, staying out of the race becomes a riskier gamble for C^1. The direct effect dominates, therefore his incentives to run are always decreasing in ψ. However, due to the indirect effect the decrease is at a slower rate in the open seat election. As a consequence, C^2's incumbency advantage is increasing in the probability of a negative shock.

Isolating the Information Channel

In the baseline model exogenous shocks influence politicians' expected utility from office via two channels: legacy (i.e., the expected value of holding office today, which here is assumed to be always lower in times of crisis) and information (i.e., the informativeness of the governance outcome, which in turn influences politician's future electoral chances). When we assume that politicians only live for two electoral cycles, both channels are necessary to generate the inefficiency documented in Proposition 1: if $\gamma = 0$ all potential candidates always choose to run for office in equilibrium. Since the value of holding office is the same in both periods, a politician would in fact never give up office today in order to increase his electoral chances tomorrow. Suppose instead that we allow players to consider a longer time horizon. Would adverse selection emerge in equilibrium even if we assume that ω influences politicians' expected utility only via the information channel (i.e., $\gamma = 0$)?

Answering this question is especially relevant in light of the results in Ashworth et al. (2017). As discussed in the robustness section, the authors show that governance outcomes are more informative (resp. less informative) following a crisis, if crises amplify (resp. mute) the effect of type. The more competence matters, the more the voter learns upon observing the incumbent's performance. As such, if the information channel alone is enough to generate the adverse selection documented in the baseline model, the key inefficiency result presented in this paper holds irrespective of whether we assume that competence matters most in times of crisis or during periods of 'business as usual'. If crises mute the effect of the office holder's type, then the voter benefits the most from a competent politician during normal times. However, this is also when outcomes are most informative. As a consequence, the politician who is most likely to be competent experiences fear of failure and has incentives to stay out of the race, running for office only during periods of crisis.

In what follows, I introduce an amended version of the model, in which politicians live for more than two periods, and the value from holding office is not a function of their performance. I will show that, if politicians are sufficiently patient, the adverse selection documented in the baseline model continues to emerge in equilibrium.

The Infinite Horizon Model

Consider a game that lasts for infinitely many periods, $t \in \{1, 2, ..., \infty\}$. At the beginning of the game each party $P \in \{1, 2\}$ randomly draws a potential candidate from the pool of its members, containing a proportion q_P of good types. Let $q_1 > q_2$. In each period, each potential candidate decides whether to run for office or not. The voter then makes her electoral decision. Office holders are subject to a two-terms limit. When an incumbent leaves office — whether because he hits the term limit, decides to stand down, or is outvoted — he cannot re-enter the pool of candidates. His party then draws a replacement (potential) candidate from the same pool. Notice that all politicians belonging to the same party are ex-ante identical.[7] This allows me to consider, in the equilibrium analysis, a generic potential candidate from Party 1 and a generic potential candidate from Party 2. As in the baseline model, when party P is unable to

[7] There is a slight technical difficulty associated with the fact that the pool depletes over time. To bypass this problem, I assume that whenever a party draws a new potential candidate, another politician with the same true type is born into the pool.

field a viable candidate it resorts to the reserve candidate R^P, that is known to be a bad type with certainty.

In each period the country experiences either a normal situation or a crisis, $\omega_t \in \{S, N\}$. Players share common prior beliefs that $prob(\omega_t = S) = \bar{p}$, with ω_t i.i.d. in each period. At the beginning of each period players observe public signal $\chi \in \{S, N\}$. For purposes of simplicity, I will assume that $prob(\chi_t = S|\omega_t = S) = prob(\chi_t = N|\omega_t = N) = 1 - \epsilon$, where ϵ takes an arbitrarily small value. In other words, the signal is (almost) perfectly informative. Notice that ϵ is assumed to be strictly larger than 0 to ensure that the voter is never indifferent between candidates of different expected quality. The production function for the governance outcomes is exactly as in the baseline model.

Politicians care exclusively about the material rents from office $K > 0$, and discount future payoffs by a common factor $\delta \in (0, 1)$. A politician's payoff when out of office is normalized to 0. The voter cares about governance outcomes, and I assume that she fully discounts the future (i.e., she maximises per-period payoff). This ensures that, in each period, the candidate who is most likely to be competent wins the election irrespective of incumbency status. This is not necessarily true in equilibrium with a forward looking voter. When choosing between a term limited incumbent and a challenger that is less likely to be competent but can run again in the following period, a forward looking voter would under some conditions elect the challenger. This is because the term limit would otherwise prevent her from efficiently using all the available information when making her electoral decision in the next period.

Finally, as in the two-period version, I assume that $\mu_{t,2}(I, S, g) > q_1$, where $\mu_{t,2}(I, \omega_{t-1}, o_{t-1})$ is the posterior probability that an incumbent from Party 2 is a good type given the previous period state of the world and governance outcome.

Analysis

The aim of this section is to verify that, under some conditions, the adverse selection documented in Proposition 1 emerges in equilibrium.[8] In this model, the problem that politicians face is to choose the right time to enter the electoral arena, so as to maximize the chances of remaining in office for two consecutive period. As such, (given $\delta < 1$) they may face a trade off between getting to office today, and waiting for a better time in order to maximize their re-election chances. [9]

Consider first a randomly drawn potential candidate from Party 1. This politician is ex-ante more likely be competent than any randomly drawn challenger from the other party. As such, he is always guaranteed re-election for a second term if he gets to office during normal times, when no new information is generated about his type. His incentives are therefore exactly as in the baseline model. He is always willing to run under $\omega_t = N$, but may decide to stay out of the race during periods of crisis in order to preserve his electoral capital and maximise the probability of getting to office when re-election is more likely. Straightforwardly, the higher the probability of being competent q_1, the stronger the incentives to run irrespective of the state of the world.

Interestingly, the opposite holds for a potential candidate from Party 2. As in the baseline model, this politician has incentives to gamble on his own success. Irrespective of how likely he is to fail, he is therefore always willing to run during times of crisis. Perhaps more surprisingly, if he is sufficiently likely to be a good type, a potential candidate from Party 2 may instead want to stay out of the race under normal times. Recall that governance outcomes are uninformative under $\omega_t = N$. Therefore, an incumbent from Party 2 would only be re-elected if his potential challenger decides to sit the election out. Conversely, a negative shock potentially allows the ex-ante disadvantaged incumbent to prove himself, thereby increasing the probability that he wins re-election even if the challenger decides to run. As such, politicians from Party 2 maximise the probability of being elected for two consecutive terms if they get to office

[8] In future iterations of the paper I will complete the analysis by characterising the Markov perfect equilibria of the infinite horizon game for all parameter values.

[9] Recall that the two-term limit implies that all incumbents will always run for re-election.

during challenging times. This, in turn, generates incentives to stay out of the race during normal ones.[10] Interestingly, as mentioned above, these incentives are stronger the higher the probability of being competent. When q_2 is high, a randomly drawn politician from Party 2 that gets elected during challenging times is very likely to survive to a second term. The opportunity cost of getting to office during normal times is too high, and the politician would rather wait for a period of crisis.

The above discussion highlights that the incentives that arise in this model are similar to those emerging in the baseline. The next proposition establishes that the equilibrium results are as well:

Proposition 3. *There exist unique $\widehat{q_2}$, $\widehat{\beta}$, and $\widehat{\delta}$ such that, if*

(i) A randomly drawn potential candidate from Party 2 is sufficiently likely to be a bad type

$0 \leq q_2 < \widehat{q_2}$

(ii) The probability that a bad type delivers a good outcome under a crisis is sufficiently low

$0 \leq \beta < \widehat{\beta}$

(iii) The politicians' discount factor is sufficiently high

$\widehat{\delta} < \delta < 1$

then, the game has an equilibrium in which any potential candidate drawn from Party 2 runs under both states of the world, whereas viable candidates drawn from Party 1 only run during normal times. During periods of crisis, Party 1 resorts to the reserve candidate R^1.

Notice that the qualitative conditions are in line with those in Proposition 1.[11] However, in contrast with the results of the baseline model, adverse selection can emerge in equilibrium for any value of q_1. For a sufficiently high discount factor, potential candidates from Party 1 choose to stay out of the race in times of crisis even if the probability of being competent is arbitrarily close to 1.

[10] Recall that I assume that when an incumbent is ousted he can never re-enter the pool of candidates.

[11] It is important to highlight that, following from the discussion above, conditions (i) and (ii) are necessary both to ensure that politicians from Party 1 choose to stay out under $\omega_t = S$ *and* that politicians from Party 2 are willing to run under $\omega_t = N$.

Conclusion

Do the right candidates choose to run at the right time? I have addressed this question by analyzing a model of repeated elections, in which potential candidates are career politicians that differ in the probability of being a competent type. The key feature of the model is that, in each period, the country faces either a normal situation or a crisis. A crisis amplifies both the importance of the office holder's competence, and the informativeness of governance outcomes. I have shown that, in a world with these features, electoral accountability may have the perverse consequence of discouraging good candidates from running precisely when the voter needs them the most. The politician who is most likely to be competent has the most to lose from information. As a consequence, if a crisis is likely, he experiences fear of failure: under some conditions, he chooses to stay out of the race so as to preserve his electoral capital for the future.

References

Ashworth, Scott, and Ethan Bueno De Mesquita. 2014. "Is voter competence good for voters?: Information, rationality, and democratic performance." American Political Science Review 108(3): 565-587.

Ashworth, Scott, Ethan Bueno de Mesquita, and Amanda Friedenberg. 2017. "Accountability and information in elections." American Economic Journal: Microeconomics 9(2): 95-138.

Ashworth, Scott, Ethan Bueno de Mesquita, and Amanda Friedenberg. 2018. "Learning about voter rationality." American Journal of Political Science 62(1): 37-54.

Besley, Timothy. 2005. "Political selection." Journal of Economic perspectives 19(3): 43-60.

Besley, Timothy, and Stephen Coate. 1997. "An economic model of representative democracy." The Quarterly Journal of Economics 112(1): 85-114.

Besley, Timothy, Jose G. Montalvo, and Marta ReynalQuerol. 2011. "Do educated leaders matter?." The Economic Journal 121(554): F205-227.

Brollo, Fernanda, Tommaso Nannicini, Roberto Perotti and Guido Tabellini. 2013. "The political resource curse." American Economic Review. 103(5): 1759-96.

Callander, Steven. 2008. "Political motivations." The Review of Economic Studies 75(3): 671-697.

Caselli, Francesco, and Massimo Morelli. 2004. "Bad politicians." Journal of Public Economics 88(3): 759-782.

Cox, Gary W and Jonathan N. Katz. 1996. "Why Did the Incumbency Advantage in US House Elections Grow?" American Journal of Political Science 40(2):478497.

Dal B, Ernesto, Pedro Dal B, and Rafael Di Tella. 2006. "Plata o Plomo?: Bribe and Punishment in a Theory of Political Influence." American Political Science Review 100(1): 41-53.

Dewan, Torun, and Rafael Hortala-Vallve. 2017. "Electoral competition, control and learning." British Journal of Political Science. 1-17.

Fedele, Alessandro, and Paolo Naticchioni. 2016. "Moonlighting politicians: motivation matters!." German Economic Review 17(2): 127-156.

Jones, Benjamin F., and Benjamin A. Olken. 2005. "Do leaders matter? National leadership and growth since World War II." The Quarterly Journal of Economics 120(3): 835-864.

Levitt, Steven D. and Catherine D. Wolfram. 1997. "Decomposing the Sources of Incumbency Advantage in the US House" Legislative Studies Quarterly 22(1):4560.

Mattozzi, Andrea, and Antonio Merlo. 2008. "Political careers or career politicians?." Journal of Public Economics 92(3): 597-608.

Messner, Matthias, and Mattias K. Polborn. 2004."Paying politicians." Journal of Public Economics 88(12): 2423-2445.

Osborne, Martin J., and Al Slivinski. 1996."A model of political competition with citizen-candidates." The Quarterly Journal of Economics 111(1): 65-96.

Strulovici, Bruno. 2009. "Learning while voting: Determinants of collective experimentation." Econometrica 78(3): 933-971.

Appendix A: Proofs

Proposition 1: *In equilibrium, both C^1 and C^2 always run for office under $\chi_1 = N$.*

Consider instead $\chi_1 = S$. Suppose that the following conditions are satisfied:

(i) The public signal is sufficiently accurate ($\psi > \underline{\psi}$)

(ii) C^2 is sufficiently unlikely to be a good type ($q_2 < \bar{q}_2$)

(iii) A bad type is sufficiently unlikely to deliver a good outcome under a crisis ($\beta < \bar{\beta}$)

Then, there exists an interval $[q_2, \bar{q}]$ such that when $q_1 \in [q_2, \bar{q}]$, C^1 chooses to stay out and Party 1 resorts to the reserve candidate R^1. Instead, C^2 always chooses to enter the race.

Proof. In the main body I have provided the proof that both candidates always choose to enter the race under $\chi_1 = N$, and that C^2 is always willing to run even under $\chi_1 = S$. Consider instead C^1's incentives under $\chi_1 = S$. Let $p_1 = prob(\omega_1 = S|\chi_1 = S) = \frac{\psi\bar{p}}{\psi\bar{p}+(1-\psi)(1-\bar{p})}$. C^1's expected utility from running in the first period is:

$$K + q_1[2\gamma + K] + (1 - q_1)[1 - p_1(1 - \beta)][\gamma + K + \gamma(1 - \bar{p}(1 - \beta)] \tag{1}$$

C^1's expected utility from staying home instead is:

$$[K + \gamma(q_1 + (1 - q_1)(1 - \bar{p}(1 - \beta))][1 - p_1 + p_1(1 - q_2)(1 - \beta)] \tag{2}$$

Thus, C^1 chooses not to run in period 1 if and only if the following condition is satisfied:

$$\begin{aligned}&[K + \gamma(q_1 + (1 - q_1)(1 - \bar{p}(1 - \beta))][1 - p_1 + p_1(1 - q_2)(1 - \beta)] > \\ &K + q_1[2\gamma + K] + (1 - q_1)[1 - p_1(1 - \beta)][\gamma + K + \gamma(1 - \bar{p}(1 - \beta)]\end{aligned} \tag{3}$$

Which reduces to:

$$q_1 < 1 - \frac{(\gamma + K)(1 + q_2 p_1(1-\beta) + \beta p_1)}{p_1(1-\beta)[2\gamma + K - \gamma\bar{p}(1 - 2\beta - q_2(1-\beta))]} = \overline{q_1} \tag{4}$$

Given $q_1 > q_2$, the above requires:

$$(1-q_2)p_1(1-\beta)(2\gamma+K-\gamma\bar{p}(1-2\beta-q_2(1-\beta)))-(\gamma+K)(1+q_2p_1(1-\beta)+\beta p_1) > 0 \tag{5}$$

The condition establishes an upper bound $q_2 < \overline{q_2}$, and must always be satisfied at $q_2 = 0$. This requires:

$$p_1[(1-\beta)(2\gamma + K - \gamma\bar{p}(1-2\beta)) - \beta(\gamma + K)] - \gamma - K > 0 \tag{6}$$

This reduces to:

$$p_1 > \frac{\gamma + K}{(1-\beta)[2\gamma + K - \gamma\bar{p}(1-2\beta)] - \beta(\gamma + K)} = \underline{p_1} \tag{7}$$

Substituting $p_1 = \frac{\psi\bar{p}}{\psi\bar{p}+(1-\psi)(1-\bar{p})}$, the above establishes a lower bound $\psi > \underline{\psi}$ and must always be satisfied at $\psi = 1$. This requires:

$$\frac{\gamma + K}{(1-\beta)[2\gamma + K - \gamma\bar{p}(1-2\beta)] - \beta(\gamma + K)} < 1 \tag{8}$$

The above establishes an upper bound $\beta < \overline{\beta}$ (and it is always satisfied at $\beta = 0$). □

Proposition 2: *Incumbency status has no effect on C^2's electoral chances under $\chi_1 = N$. Suppose instead that $\chi_1 = S$. Then, C^2 experiences an incumbency disadvantage whenever $q_r > q_1$, and an advantage whenever $q_r < q_1$. In both cases, the effect of incumbency is increasing in the signal's accuracy ψ.*

Proof. The first point follows straightforwardly from the proof of Proposition 1, and so does the existence of an incumbency disadvantage under $q_r > q_1$. Suppose instead that $q_r < q_1$. C^1's utility from running in period 1 is exactly as in the baseline:

$$K + q_1[2\gamma + R] + (1 - q_1)[1 - p_1(1 - \beta)][\gamma + K + \gamma(1 - \bar{p}(1 - \beta)] \tag{9}$$

Conversely, if C^1 chooses not to run he will always win the second period election. His expected utility is therefore:

$$K + \gamma(q_1 + (1 - q_1)(1 - \bar{p}(1 - \beta)) \tag{10}$$

Thus, C^1 chooses not to run in period 1 if and only if the following condition is satisfied:

$$K+\gamma(q_1+(1-q_1)(1-\bar{p}(1-\beta)) > K+q_1[2\gamma+K]+(1-q_1)[1-p_1(1-\beta)][\gamma+K+\gamma(1-\bar{p}(1-\beta)] \tag{11}$$

Which reduces to:

$$q_1 < 1 - \frac{(\gamma + K)}{p_1(1 - \beta)[\gamma + K + \gamma(1 - \bar{p}(1 - \beta)]} \tag{12}$$

C^2's incumbency advantage is therefore:

$$1 - \frac{(\gamma + K)}{p_1(1 - \beta)[\gamma + K + \gamma(1 - \bar{p}(1 - \beta)]} - [1 - \frac{(\gamma + K)(1 + q_2 p_1(1 - \beta) + \beta p_1)}{p_1(1 - \beta)(2\gamma + K - \gamma\bar{p}(1 - 2\beta - q_2(1 - \beta)))}] > 0 \tag{13}$$

Substituting $p_1 = \frac{\psi\bar{p}}{\psi\bar{p}+(1-\psi)(1-\bar{p}}$, it is easy to verify that the advantage is increasing in ψ.

□

Proposition 3:*There exist unique* $\widehat{q_2}$, $\widehat{\beta}$, *and* $\widehat{\delta}$ *such that, if*

(i) A randomly drawn potential candidate from Party 2 is sufficiently likely to be a bad type $0 \leq q_2 < \widehat{q_2}$

(ii) The probability that a bad type delivers a good outcome under a crisis is sufficiently low $0 \leq \beta < \widehat{\beta}$

(iii) The politicians' discount factor is sufficiently high

$\widehat{\delta} < \delta < 1$

then, the game has an equilibrium in which any potential candidate drawn from Party 2 runs under both states of the world, whereas viable candidates drawn from Party 1 only run during normal times. During periods of crisis, Party 1 resorts to the reserve candidate R^1.

Proof. Denote as $U_P^e(H_t, \chi_t, e_t)$ the expected discounted payoff of a non-incumbent potential candidate from party $P \in \{1, 2\}$ if he chooses to enter the race at time t. $H_t \in \{1, 2\}$ indicates the identity of the potential candidate with the highest probability of being a good type. $e_t \in \{I, \emptyset\}$, where $e_t = \emptyset$ denotes that the race at time t is open seat and $e_t = I$ that the incumbent from the other party is running for re-election. $U_P^o(H_t, \chi_t, e_t)$ denotes the expected discounted payoff of a non-incumbent potential candidate from party P if he chooses to stay out of the race at time t.

As discussed in the main body, non-incumbent potential candidates from Party 1 are always willing to run under $\chi_t = N$, and non-incumbent potential candidates from Party 2 are always willing to run under $\chi_t = S$. Further, all incumbents are always willing to run for re-election.

Consider instead a potential candidate from Party 2 under $\chi_t = N$. In the conjectured equilibrium, he is always indifferent between running for office and staying home if the election is open seat, since he would lose with probability 1. Consider his entry decision when an incumbent from Party 1 is up for re-election, and performed poorly in the previous period. In the conjectured equilibrium, his expected discounted payoff is:

$$U_2^e(2, N, I) = K + \delta K \bar{p} \tag{14}$$

Since he would only win re-election for a second term if the public signal indicates a crisis and therefore the (new) potential candidate from Party 1 chooses to stay home.

His expected discounted payoff from a deviation would instead be:

$$\delta(\bar{p}U_2^e(1,S,\emptyset) + (1-\bar{p})U_2^e(1,N,\emptyset)) \tag{15}$$

Where

$$U_2^e(1,S,\emptyset) = K + \delta K[\bar{p} + (1-\bar{p})(q_2 + (1-q_2)\beta)] \tag{16}$$

And

$$U_2^e(1,N,\emptyset) = U_2^o(1,N,\emptyset) = \delta^2(\bar{p}U_2^e(1,S,\emptyset) + (1-\bar{p})U_2^e(1,N,\emptyset)) \tag{17}$$

Remember that the public signal is (almost) perfectly informative (since I assume $prob(\chi_t = S|\omega_t = S) = prob(\chi_t = N|\omega_t = N) = 1-\epsilon$), and I can therefore ignore the arbitrarily small probability that a crisis arises after a signal $\chi_t = N$.

Solving for $U_2^e(1,N,\emptyset)$ we obtain that the deviation is not profitable if and only if the following condition is satisfied:

$$K + \delta K\bar{p} > \delta K\bar{p}\frac{(1+\delta(\bar{p} + (1-\bar{p})(q_2 + (1-q_2)\beta))}{1-\delta^2(1-\bar{p})} \tag{18}$$

Rearranging we obtain:

$$q_2 < \frac{1-\delta^2(\bar{p}^2 + (1-\bar{p})(1+\delta\bar{p}))}{\delta^2\bar{p}(1-\bar{p})(1-\beta)} - \frac{\beta}{1-\beta} \tag{19}$$

Since, $q_2 > 0$ the above requires:

$$\beta < \frac{1-\delta^2(\bar{p}^2 + (1-\bar{p})(1+\delta\bar{p}))}{\delta^2\bar{p}(1-\bar{p})} \tag{20}$$

Consider now a non-incumbent potential candidate from party 1 under $\chi_t = C$. Intuitively, his incentives to run are stronger when a term limited incumbent is up for re-election (as compared to an open seat election). As such, it is sufficient to show that the equilibrium is robust to a deviation in this case.

Considering the case in which $H_t = 1$, Party 1's potential candidate expected discounted payoff in the conjectured equilibrium is:

$$U_1^o(1,S,2) = \delta((1-\bar{p})U_1^e(1,N,\emptyset) + \bar{p}U_1^o(1,S,\emptyset)) \tag{21}$$

Where

$$U_1^e(1,N,\emptyset) = K(1+\delta) \tag{22}$$

And

$$\begin{aligned} U_1^o(1,C,\emptyset) = \delta[1-(q_2+(1-q_2)\beta)](\bar{p}U_1^o(1,C,2) + (1-\bar{p})U_1^e(1,N,2) \\ +\delta[q_2+(1-q_2)\beta)]\bar{p}(U_1^o(2,C,2) + (1-\bar{p})U_1^e(2,N,2)) \end{aligned} \tag{23}$$

With $U_1^e(2,N,2) = U_1^o(2,N,2) = \delta(\bar{p}U_1^o(1,S,\emptyset)+(1-\bar{p})U_1^e(1,N,\emptyset))$ and $U_1^e(1,N,2) = U_1^e(1,N,\emptyset)$.

His expected discounted payoff from a deviation is instead:

$$K + \delta K(q_1 + (1-q_1)\beta) \tag{24}$$

Solving for $U_1^o(1,S,\emptyset)$ and rearranging we obtain that the deviation is not profitable if and only if the following condition is satisfied:

$$(K+\delta K)\delta(1-\bar{p})\frac{(1-q_2)(1-\beta)+\delta(\bar{p}+(1-\bar{p})(q_2+(1-q_2)\beta)}{1-\delta^2\bar{p}[\bar{p}+(1-\bar{p})(q_2+(1-q_2)\beta)]} > K+\delta K(q_1+(1-q_1)\beta) \tag{25}$$

Rearranging we obtain:

$$q_2 < \frac{(1+\delta\bar{p})[\delta(1+\delta)(1-\bar{p}) - (1-\delta\bar{p})(1+\delta(q_1+(1-q_1)\beta))]}{\delta(1-\bar{p})(1-\beta)[1-\delta^2(1-\bar{p}(1-q_1)(1-\beta))]} - \frac{\beta}{1-\beta} \tag{26}$$

This requires

$$\frac{(1+\delta\bar{p})[\delta(1+\delta)(1-\bar{p}) - (1-\delta\bar{p})(1+\delta(q_1+(1-q_1)\beta))]}{\delta(1-\bar{p})(1-\beta)[1-\delta^2(1-\bar{p}(1-q_1)(1-\beta))]} - \frac{\beta}{1-\beta} > 0 \tag{27}$$

The above condition establishes an upper bound $\beta < \tilde{\beta}$. $\tilde{\beta} > 0$ requires

$$(1 + \delta\bar{p})[\delta(1 + \delta)(1 - \bar{p}) - (1 - \delta\bar{p})(1 + \delta q_1)] > 0 \tag{28}$$

The LHS is increasing in δ, fails at $\delta = 0$ and is always satisfied at $\delta = 1$. The condition therefore establishes a lower bound $\delta > \widehat{\delta}$

Thus, the conjectured equilibrium exists if and only if the following conditions are satisfied:

- $0 < q_2 < \widehat{q_2} = \min\{\frac{1-\delta^2(\bar{p}^2+(1-\bar{p})(1+\delta\bar{p}))}{\delta^2\bar{p}(1-\bar{p})(1-\beta)} - \frac{\beta}{1-\beta}, \frac{(1+\delta\bar{p})[\delta(1+\delta)(1-\bar{p})-(1-\delta\bar{p})(1+\delta(q_1+(1-q_1)\beta))]}{\delta(1-\bar{p})(1-\beta)[1-\delta^2(1-\bar{p}(1-q_1)(1-\beta))]} - \frac{\beta}{1-\beta}\}$
- $\beta < \widehat{\beta} = \min\{\tilde{\beta}, \frac{1-\delta^2(\bar{p}^2+(1-\bar{p})(1+\delta\bar{p}))}{\delta^2\bar{p}(1-\bar{p})}\}$
- $\delta > \widehat{\delta}$

□

Appendix B: Robustness

In this section I formally analyse the variants of the baseline model introduced in the Discussion and Robustness section.

Governance outcomes directly influence politicians' payoffs

Consider an amended version of the baseline mode in which politicians' payoffs are as follows:

- $K + \mathbb{I}_g\gamma - (1-\mathbb{I}_g)\lambda$ when in office
- $-(1-\mathbb{I}_g)\lambda$ when not in office

Where $\mathbb{I}_g$ is a binary indicator taking value 1 if $o_t = g$ and 0 otherwise.

In equilibrium, C^1 chooses not to run in the first period if and only if the following condition is satisfied:

$$\begin{aligned} p_1(1-\beta)(1-q_2)(\gamma(1-(1-\beta)(1-q_1)\bar{p}) - (1-\beta)(1-q_1)\lambda\bar{p} - \lambda + K) \\ -p_1\beta(1-q_2)(1-\beta)\bar{p}\lambda \\ +(1-p_1)(\gamma(1-(1-\beta)(1-q_1)\bar{p}) - (1-\beta)(1-q_1)\lambda\bar{p} + K) > \\ K + q_1(2\gamma + K) \\ +(1-q_1)(1-(1-\beta)p_1)(\gamma(1-(1-\beta)\bar{p}) - (1-\beta)\lambda\bar{p} + \gamma + K) \\ -p_1(1-\beta)(1-q_1)((1-\beta)(1-q_2)\lambda\bar{p} + \lambda) \end{aligned} \tag{29}$$

This reduces to:

$$q_1 < 1 - \frac{(\gamma+K)(1+q_2p_1(1-\beta)+\beta p_1) + \lambda(1-\beta)(1-q_2)p_1(1+\beta\bar{p})}{p_1(1-\beta)[2\gamma + K - \gamma\bar{p}(1-2\beta-q_2(1-\beta)) + \lambda(1+\bar{p}\beta)]} = \overline{q_1}_\lambda \tag{30}$$

Given $q_1 > q_2$, the above requires:

$$(1-q_2)p_1(1-\beta)[2\gamma + K - \gamma\bar{p}(1-2\beta - q_2(1-\beta)) + \lambda(1+\bar{p}\beta)] \\ -(\gamma+K)(1+q_2p_1(1-\beta)+\beta p_1) - \lambda(1-\beta)(1-q_2)p_1(1+\beta\bar{p}) > 0 \tag{31}$$

The LHS is decreasing in q_2, therefore the condition establishes an upper bound $q_2 < \overline{q_2}_\lambda$ and must be satisfied at $q_2 = 0$.

$$p_1(1-\beta)(2\gamma+K-\gamma\bar{p}(1-2\beta)+\lambda(1+\bar{p}\beta))-[(\gamma+K)(1+\beta p_1)+\lambda p_1(1-\beta)(1+\beta\bar{p})] > 0 \tag{32}$$

The inequality can only be satisfied if the LHS is increasing in p_1. Substituting $p_1 = \frac{\psi\bar{p}}{\psi\bar{p}+(1-\psi)(1-\bar{p})}$, the above establishes a lower bound $\psi > \underline{\psi}_\lambda$ and must always be satisfied at $\psi = 1$:

$$(1-\beta)[2\gamma+K-\gamma\bar{p}(1-2\beta)+\lambda(1+\bar{p}\beta)]-[(\gamma+K)(1+\beta)+\lambda(1-\beta)(1+\beta\bar{p})] > 0 \tag{33}$$

The above is concave in β, and always at $\beta = 0$, therefore the condition establishes an upper bound $\beta < \overline{\beta}_\lambda$.

A bad governance outcome increases the probability of a crisis arising in the future

Suppose that politicians only care about their own performance in office, and consider an amended version of the baseline model where the probability of a negative shock in the second period is a function of the first period governance outcome:

- $prob(\omega_2 = C|o_1 = g) = \bar{p}$
- $prob(\omega_2 = C|o_1 = b) = \alpha\bar{p}$, where $\alpha \in (1, \frac{1}{\bar{p}})$

C^1 will choose not to run in period 1 if and only if the following condition is satisfied:

$$\begin{aligned}[K+\gamma(q_1+(1-q_1)(1-\bar{p}(1-\beta))](1-p_1) \\ +p_1(1-q_2)(1-\beta)][K+\gamma(q_1+(1-q_1)(1-\alpha\bar{p}(1-\beta))] > \\ K+q_1[2\gamma+K]+(1-q_1)[1-p_1(1-\beta)][\gamma+K+\gamma(1-\bar{p}(1-\beta)]\end{aligned} \tag{34}$$

This reduces to:

$$q_1 < 1-\frac{(\gamma+K)(1+q_2p_1(1-\beta)+\beta p_1)}{p_1(1-\beta)[2\gamma+K-\gamma\bar{p}(\alpha(1-q_2)(1-\beta)-\beta)]} = \overline{q_1}_\alpha \tag{35}$$

Given $q_1 > q_2$, the above requires:

$$1-q_2-\frac{(\gamma+K)(1+q_2p_1(1-\beta)+\beta p_1)}{p_1(1-\beta)[2\gamma+K-\gamma\bar{p}(\alpha(1-q_2)(1-\beta)-\beta]} > 0 \tag{36}$$

Substituting $p_1 = \frac{\psi\bar{p}}{\psi\bar{p}+(1-\psi)(1-\bar{p})}$, the above establishes a lower bound $\psi > \underline{\psi}_\alpha$ and must always be satisfied at $\psi = 1$:

$$(1-q_2)(1-\beta)[2\gamma+K-\gamma\bar{p}(\alpha(1-q_2)(1-\beta)-\beta)]-(\gamma+K)(1+q_2(1-\beta)+\beta) > 0 \tag{37}$$

The LHS is decreasing in q_2, therefore it establishes an upper bound $q_2 < \overline{q_2}_\alpha$ and must always be satisfied at $\overline{q_2} = 0$:

$$(1-\beta)[2\gamma+K-\gamma\bar{p}(\alpha(1-q_2)(1-\beta)-\beta)]-(\gamma+K)(1+\beta) > 0 \tag{38}$$

The LHS is concave in β and always satisfied at $\beta = 0$. Thus, it establishes an upper bound $\beta < \overline{\beta}_\alpha$.

A bad governance outcome decreases the country's future resiliency

Suppose that politicians only care about their own performance in office, and the probability of a crisis in the second period is exogenous. Consider an amended version of the baseline model in which the first period governance outcome influences the probability that $o_2 = g$ if the country experiences a crisis and the office holder is a bad type:

- $prob(o_2 = g|\omega_2 = C, \theta_{I_2} = B, o_1 = g) = \beta$
- $prob(o_2 = g|\omega_2 = C, \theta_{I_2} = B, o_1 = g) = \delta\beta$, where $\delta \in [0, 1]$

C^1 chooses not to run in the first period if and only if the following condition is satisfied:

$$\begin{aligned} p_1(1-q_2)(1-\beta)[K+\gamma[q_1+(1-q_1)(1-\bar{p}(1-\delta\beta))]+ \\ (1-p_1)[K+\gamma[q_1+(1-q_1)(1-\bar{p}(1-\beta))]]> \\ K+q_1(2\gamma+K)+(1-q_1)(1-p_1(1-\beta))(K+\gamma(2-\bar{p}(1-\beta))) \end{aligned} \tag{39}$$

This reduces to:

$$q_1 < 1 - \frac{(\gamma+K)(1+q_2p_1(1-\beta)+\beta p_1)}{p_1(1-\beta)[2\gamma+K-\gamma\bar{p}(1-\beta(1+\delta)-q_2(1-\beta))]} = \overline{q_1}_\delta \tag{40}$$

Given $q_1 > q_2$, this requires:

$$(1-q_2)p_1(1-\beta)[2\gamma+K-\gamma\bar{p}(1-\beta(1+\delta)-q_2(1-\beta))]-(\gamma+K)(1+q_2p_1(1-\beta)+\beta p_1) > 0 \tag{41}$$

The above establishes an upper bound $q_2 < \overline{q_2}_\delta$. Thus, the condition must be satisfied at $q_2 = 0$. This requires:

$$p_1(1-\beta)[2\gamma+K-\gamma\bar{p}(1-\beta(1+\delta)]-(\gamma+K)(1+\beta p_1) > 0 \tag{42}$$

Substituting $p_1 = \frac{\psi\bar{p}}{\psi\bar{p}+(1-\psi)(1-\bar{p})}$, the above establishes a lower bound $\psi > \underline{\psi}_\delta$ and must always be satisfied at $p_1 = 1$:

$$(1-\beta)[2\gamma+K-\gamma\bar{p}(1-\beta(1+\delta)]-(\gamma+K)(1+\beta) > 0 \tag{43}$$

The LHS is concave in β, and it is always satisfied at $\beta = 0$. The condition therefore establishes an upper bound $\beta < \overline{\beta}_\delta$.

State-dependent legacy payoffs

Consider an amended version of the baseline model in which an office holder's legacy payoff from a good performance is higher under $\omega_t = C$:

- K if $o_t = b$
- $K + \gamma$ if $o_t = g$ and $\omega_t = N$
- $K + \nu(\omega_t)\gamma$, where $\nu(S) > 1$ and $\nu(N) = 1$

C^1 chooses not to run in the first period if and only if the following condition is satisfied:

$$\begin{aligned}
&[K + \gamma(1 - \bar{p} + \nu(S)\bar{p}(q_1 + (1 - q_1)\beta))][1 - p_1 + p_1(1 - q_2)(1 - \beta)] > \\
&K + q_1[K + \gamma(2 + (p_1 + \bar{p})(\nu(S) - 1)] + (1 - q_1)[1 - p_1][\gamma + K + \gamma(1 - \bar{p} + \nu(S)\bar{p}\beta)] \\
&+(1 - q_1)p_1\beta[\nu(S)\gamma + K + \gamma(1 - \bar{p} + \nu(S)\bar{p}\beta)]
\end{aligned} \tag{44}$$

This reduces to:

$$q_1 < \frac{p_1((1 - 2\beta - q_2(1 - \beta))(K + \gamma(1 - \bar{p}(1 - \beta\nu(S)))) + \gamma(1 - \beta\nu(S))) - (K + \gamma)}{p_1(1 - \beta)(\gamma(1 + \nu(S)) + K - \gamma\bar{p}(1 - 2\nu(S)\beta - \nu(S)q_2(1 - \beta))} = \overline{q_1}_\nu(S) \tag{45}$$

Given $q_1 > q_2$, the above requires:

$$\begin{aligned}
&p_1[(1 - 2\beta - q_2(1 - \beta))(K + \gamma(1 - \bar{p}(1 - \beta\nu(S)))) + \gamma(1 - \beta\nu(S))] - (K + \gamma) \\
&-q_2[p_1(1 - \beta)(\gamma(1 + \nu(S)) + K - \gamma\bar{p}(1 - 2\nu(S)\beta - \nu(S)q_2(1 - \beta))] > 0
\end{aligned} \tag{46}$$

The LHS is decreasing in q_2, therefore it establishes an upper bound $q_2 < \overline{q_2}_\nu$ and must always be satisfied at $q_2 = 0$:

$$p_1((1-2\beta)(K+\gamma(1-\bar{p}(1-\beta\nu(S))))+\gamma(1-\beta\nu(S)))-(K+\gamma)>0 \tag{47}$$

Substituting $p_1 = \frac{\psi\bar{p}}{\psi\bar{p}+(1-\psi)(1-\bar{p})}$, the above establishes a lower bound $\psi > \underline{\psi}_\nu$ and must always be satisfied at $\psi = 1$:

$$(1-2\beta)(K+\gamma(1-\bar{p}(1-\beta\nu))+\gamma(1-\beta\nu(S)))-(K+\gamma)>0 \tag{48}$$

The LHS is concave in β, and it is always satisfied at $\beta = 0$ The condition therefore establishes an upper bound $\beta < \overline{\beta}_\nu$.