The London School of Economics and Political Science

Essays on Dynamic Political Economy

Ana McDowall

A thesis submitted to the Department of Economics of the London School of Economics and Political Science for the degree of Doctor of Philosophy
London, May 2015

To Valle and Peter

Declaration

I certify that the thesis I have presented for examination for the 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).

I warrant that this authorization does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of 52,020 words.

Abstract

Many of the problems faced by the agents to whom we delegate public decisions involve dynamic considerations. Whether it is a public official whose reputation is on the line, or a politician who faces repeated elections or who worries about how successors will alter his legacy, decisions today are made with attention to what consequences may come tomorrow. In this dissertation, I discuss dynamic problems in political economy. The common thread to all three essays will be that decision makers face dynamic incentives to protect their own interest because their policy choice today affects tomorrow's decision environment and, as a result, how other agents (the public, or rival parties) will behave further down the line.

The first chapter looks at how alternation in power of (pro-rich and pro-poor) partisan political groups affects incentives to implement short and long run redistributive policies. I identify a powerful incentive for both groups to make income-equalising investments as a form of insurance against takeover by the opposing group.

The second chapter studies the pressures on an appointed regulator of some risky activity in society, who cares about protecting his reputation. The regulator is held to account by the general public, which use a heuristic approach to estimating risks and assessing the performance of the regulator. The public official will trade off trying to align his policy record with current beliefs and managing how beliefs change through observing the risky activity.

In the third chapter, I consider parties that compete over a one-dimensional ideological policy space but face uncertain electoral outcomes. I find that if elections are systematically biased against the incumbent, this leads to higher ideological polarisation than a bias of the same size in favour of the incumbent, because it lowers the value of winning. I endogenise incumbency biases with a model where voters learn about candidates' competence.

Acknowledgments

The work contained in this thesis has benefited from the support of many people. I would like to thank, first and foremost, my supervisor Gilat Levy for her unfailing encouragement, and for her invaluable contributions to my thinking as I developed the ideas for this thesis. From the first day to the last, she has supported me in countless ways; she has helped me to understand my own work better, and to not lose momentum. I am also indebted to my advisor Michele Piccione for his curiosity, which pushed me to think harder about problems; it has been a privilege to share many discussions of work, music and many topics with him.

I would like to thank all the Theory faculty at the LSE, for being a generally brilliant group of thinkers and for always having time for and interest in their students' projects and ideas. Also at the LSE, I thank Tim Besley, Rafael Hortala-Vallve and Ethan Ilzetzki for their helpful comments. Finally, I am very grateful to Daniel Diermeier, Sandeep Baliga and Eddie Dekel for their time and thoughts during my brief passage through Kellogg School of Management at Northwestern.

I have been lucky enough to share the PhD adventure with some fantastic economists-in-the-making. I thank Laura Derksen, for being a patient, smart, and generous friend and flatmate. Sarah Auster, Patrick Blanchenay, Lena Koerber, Samuel Marden, Markus Riegler, Lisa Windsteiger and Kelly Zhang have all impressed me with their talent and their enthusiasm for the subject, and have been fun and kind colleagues. From my time in the US, I thank Enrico Berkes, José Antonio Espín Sánchez, René Leal, Anand Krishnamurthy and Paul Mohnen for being my economics family in Evanston.

I thank the friends that have been there all along, Aoife Bharucha, Robbie Campbell, Alejandra Climent, Esther Corral, Beatriz López-Fanjul, Marta Medina and Paula Millán; as well as the ones I've made along the way, Kaylyn Fraser, James Lyon and Julia Poncela. They have all been there to celebrate the highs and to weather the lows. I also thank Kilian Huber for his humour and calmness in the final stretch.

Finally, I am eternally grateful to my family, my incredible parents Peter and Valle and sister Clare, who are my biggest supporters and the source of all my strength.

Financial support from the Economic and Social Research Council is most gratefully acknowledged.

Contents

List of Figures

Chapter 1

Public Investment and the Dynamics of Inequality

When there is turnover in government between partisan groups that defend the interests of the rich and the poor, a conflict exists over policies to redistribute in any given period. Governments can also take measures to affect the future distribution of income, for example through investments in human capital. I provide a formal model in which static and dynamic policies for redistribution interact, in the presence of political alternation. Investments that alter inequality are used strategically by governing parties to mitigate future redistributive conflict. Specifically, parties on either side choose to maintain lower inequality in gross incomes when they are less likely to be in office in future.

1.1 Introduction

Governments have the ability to influence the distribution of welfare by taxing citizens, financing public goods and implementing transfers. In addition to this, they have the ability to influence the future pre-tax distribution of income, and therefore welfare, by making durable investments in public capital. These policy decisions are made in the context of electoral competition which provides constraints on the policies that governments will choose, and, crucially, creates strategic incentives for policy selection. A party in power will try to influence future conditions to its advantage, whether this means increasing its probability of holding onto power, or constraining policy choices for a government with different preferences. In this paper the fundamental conflict between parties will concern preferences over redistribution.

Previous work has extensively studied this redistributive power of political institutions. Some have studied the effect of inequality on the evolution of the political regime; Acemoglu and Robinson 2001 point out that when inequality becomes large, there will be increasing pressure on institutions to democratise as a mechanism to commit to more generous redistribution. Another line of research considers how redistributive policies chosen democratically affect the macroeconomy; Persson and Tabellini 1994, Bertola 1991 and Alesina and Rodrik 1991 all argue that high inequality hinders growth because the government redistributes more aggressively, creating disincentives for private investment.

My aim is to extend the analysis of redistribution policy to include two dynamic features: 1) public investments that influence the future income distribution, and 2) uncertainty about the identity of future governments. The focus of this work is on the strategic effects that arise in an environment where there is a threat of losing office to a rival party with different redistributive preferences, and the role of investment as a form of dynamic redistribution that serves to insure against political turnover.

Public investments as I define them (any policy that alters future incomes), cover a variety of government programmes. Governments can raise the productivity of groups by rolling out literacy programmes, investing in primary education, offering retraining for workers in declining industries, enabling migration to cities, or opening up credit lines for the rural sector. They can implement land reform or reallocate other productive assets. Governments can also influence future income distribution through policies that affect social mobility: introducing labour laws to prevent discrimination, improving education opportunities for individuals of deprived backgrounds, or providing health-care and nutritional support for children living in poverty. Finally, many policies have effects on both within-period and across-periods redistribution; I separate the two to display the mechanism more clearly.

Electoral uncertainty is also an important feature of any political system. First we know that most modern democracies do experience alternation between partisan groups with differing economic/redistributive agendas. In the US, between 1945 and 2010 the presidency switched between the two main parties every 8 years on average[1], so parties do not expect to hold on to power for much longer than two terms. These parties have different redistributive aims: Bartels 2009 presents evidence that on average Republicans have presided over periods of growing inequality, while Democrats have been able to slow down or even reverse the trend of growing inequality.

Electoral uncertainty occurs because voter preferences are multidimensional (they may care about politicians' charisma, competence or non-economic issues), uncertain and shifting with time. This alternation is likely to introduce distortions to decisions about dynamic redistribution, because parties will use the instruments available to them today to secure more favourable future outcomes for their group. Even in autocratic regimes, the threat of turnover exists to some extent, and those in power will consider such a prospect when making fiscal decisions.

In the model, two parties, representing the two social groups (rich and poor), compete for office. The winner is able to choose the level of (non-storable) public consumption as well as the size of investment in public capital, all financed through a proportional tax on voters. Public consumption benefits all equally, but proportional taxation means the richer group face a bigger burden for financing it. The level of inequality (ratio of incomes) determines the marginal cost of public funds for each group. Public capital affects future pre-tax incomes, and may do so differentially for the two groups, changing future inequality.

In Section 1.2, I consider a two period model with endogenous elections. Parties have no commitment to policies prior to the election and voters correctly anticipate their policy choices. Voters also value some orthogonal dimension of quality, which is modeled as a random aggregate taste shock. The median voter is poor and so the winner of the election is the preferred party of the poor types.

The two-period model delivers the basic mechanism by which static and dynamic redistribution interact. In the second period there is no investment, only public consumption. Unsurprisingly, the poorer group prefers to set a higher tax rate and fund a higher level of public consumption since it faces a lower marginal cost of public funds. The rich group prefers more modest taxation and spending. Public

[1] In fact, it changed precisely every 8 years on all occasions except two.

consumption involves static redistribution and the existence of inequality generates a redistributive conflict (a wedge between the preferred levels of public consumption of the two groups), and this conflict increases with the degree of inequality.

The choice over public investment must take into account four different effects: 1) changes to ones own private income, 2) changes to the price of public funds, 3) a strategic motive to influence the next government's choice over public consumption, in case the opposition were to take over in office, and 4) an electoral motive.

The effect of public investment on inequality will turn out to be critical in the model because, as we saw, higher inequality drives larger redistributive conflict. The first period incumbent will be concerned about insuring himself against the possibility of losing power, and when he is more preoccupied with this possibility, reducing future redistributive conflict will become more critical. This is the first comparative static result of the paper: a stronger incumbency advantage leads to higher investment if investment increases inequality, and to less investment if it reduces inequality. This is true regardless of the type of the party, both the poor and rich type will allow future inequality to be higher when they are more entrenched in office.

In fact, this strategic motive may induce positive levels of investment by the rich group even in some cases where their group is directly harmed by the investment, such as the case of purely redistributive investments. It is a fairly universal feature of modern democracies that parties of all sides support, although to varying degrees, long term poverty reduction. Investments in public education and health, or infrastructure projects to improve sanitation, housing and transport for the poorest communities, are often willingly implemented by political elites that are right-of-centre. The analysis I present suggests that along with purely altruistic motives (or reasons related to aggregate growth), there is a strong strategic incentive for the parties of the rich to do so, as long as the threat of political alternation is strong enough.

I provide two more comparative static results in the two period model. They consider how investment choices are affected by office rents and by the responsiveness of voting to policies (i.e. the variance of the aggregate taste shock). Higher office rents make parties care more about winning per se and less about providing policies that suit their group. Higher rents will drive the rich group to go the extra mile to reduce redistributive conflict, that is, invest more in income-equalising measures. In contrast, the poor group will allow for higher inequality, since higher redistributive conflict provides their party with an edge over the rich party in tomorrow's election. The other result is that when voting becomes more sensitive to policies, manipulation of future electoral prospects through investment choices is more effective. The median voter penalises the rich party more heavily for providing a less desirable platform. Since it has no commitment to policies, the rich party has a stronger incentive to lock in a lower level of redistributive conflict. The party of the poor exploits this higher sensitivity by exacerbating the conflict from higher inequality.

I extend the analysis to a fully dynamic model with repeated interaction and exogenous elections, focusing on Markov Perfect Equilibria (MPE). With repeated investment decisions, my first result is about how investment policies change as inequality evolves: policy functions are monotonic and opposite for the two groups. If one group has a higher rate of income growth with respect to public capital than the other, then this group invests less at higher initial levels of public capital, while the rival group invests more. This is driven by a price effect: the price of investment for the former group rises with public capital, while the latter faces lower prices.

The strategic effects described earlier reappear in the fully dynamic environment. I work with a linear quadratic framework which can be solved numerically for any given parameter configuration. The comparative static result for incumbency advantage holds: both parties choose lower levels of inequality in the next period when they are less likely to hold onto power.

I consider welfare and a comparison of outcomes under a utilitarian central planner and under the strategic environment with political turnover. A central planner chooses intermediate levels of public consumption compared to the choices of partisan politicians, but may invest more (or less) than both politicians. A central planner's choice of investment is high if the investment generates high growth in aggregate income, but is not concerned at all with the effect on inequality. For example, when public investment is purely redistributive, the central planner won't invest, but the politicians might, even the ones representing the group whose income will shrink with investment.

This suggests an interpretation of the strategic effect as resolving a commitment problem. The central planner faces no commitment problem and hence has no need for dynamic redistribution. In contrast, the politicians cannot commit to moderate levels of static redistribution, and will use investment to tie the hands of their opposition to some extent[2]. The rich will reduce inequality to prevent over-spending by the poor; the poor will reduce inequality to make the rich party more willing to spend.

The existence of redistributive conflict hurts both groups. Welfare may increase for both groups when there is lower incumbency advantage, even for the current

[2]This result mirrors Alesina and Tabellini 1990 in which inefficient government deficits appear in order to tie the hands of future governments who may have different preferences over public spending.

incumbent. With a weaker hold over power, parties will maintain lower levels of inequality and redistributive conflict will be lower, reducing the variance in both groups' period payoffs. Welfare may also increase when initial levels of inequality are lower, even if this means one group starting out with lower income, the reason being again a reduction in redistributive conflict.

The paper proceeds as follows. In the remainder of Section 1.1, I provide a brief discussion of related literature. Section 1.2 presents the two period model with endogenous elections. Section 1.3 discusses welfare and the benchmark of outcomes under delegation to a utilitarian central planner. Section 1.4 provides results on the features of MPE in a general infinite horizon model with exogenous elections. Since a full characterisation of the equilibrium is not possible for the most general model, Section 1.5 presents a linear-quadratic framework; I confirm the key comparative statics and discuss new results arising in the fully dynamic setting. Section 1.6 concludes.

Related Literature

There is an extensive literature that considers redistributive policies as the result of a political process. Using Meltzer and Richard 1981 as a workhorse model for the determination of redistribution in a democracy, several papers have studied how inequality affects economic growth. Alesina and Rodrik 1994, Persson and Tabellini 1994 and Bertola 1991 predict that higher inequality (measured by the gap between median and mean incomes) generates excessive transfers which disincentivise private investment[3]. The predictions were tested empirically (Persson and Tabellini 1992, Perotti 1992, or Perotti 1996), with mixed results. Although a correlation between inequality and slow growth is observed, the link between more aggressive redistribution and slower growth is not significant and is sometimes contradicted in the data. A new breed of models (Galor and Zeira 1993, Perotti 1993, Benabou 1996, Aghion and Bolton 1997) extended the formal analysis to reconcile the theory with the empirical findings.

I depart from these models in two important ways. First, I assume governments can also engage in dynamic redistribution and affect the future composition of gross earnings in society. Second, I depart from the median voter result. I assume there is electoral uncertainty and parties with different preferences over redistribution have positive probability of winning office. While expanding these elements of the

[3]These models offer small variations on the particular use of public revenues. In Persson and Tabellini 1994 they are straightforward transfers, and since taxation is costly, redistribution is wasteful. In Alesina and Rodrik 1994 revenues go toward producing a valuable public good, and in Bertola 1991 they correct an externality.

environment I simplify others. Most importantly, I abstract entirely from private investment.

This paper mainly contributes to the literature on the strategic effects of political turnover on policy outcomes. Alesina and Tabellini 1990 and Persson and Svensson 1989 propose that in the presence of alternation in government between parties with different fiscal preferences[4], government debt will be inefficiently high. Parties in office overspend today to constrain the spending power of future governments. Besley and Coate 1998 and Azzimonti 2014 consider public investment choices in the presence of political uncertainty, focusing on policy inefficiency ("political failure"). Investment choices are inefficient because tomorrow's government may have different preferences over future expenditures. In all of the above, the strategic effects on policy of a threat of replacement by a group with different preferences, is common to my work. The key difference is that in these papers, today's government can change the constraints faced by future governments (typically available resources) while in my work they are also able to influence future preferences.

Another related area of research studies the relationship between inequality and the institutional environment. Acemoglu and Robinson's theory of democratisation proposes how different levels of inequality affect what type of institution is likely to emerge in equilibrium. Extending the franchise is seen as a form of commitment to higher future redistribution by shifting the identity of the median voter. In a sense, this paper looks at the link running in the opposite direction. Instead of predicting institutional transitions for a given level of inequality, I analyse how inequality will evolve, for a given institutional setting.

In that sense, Austen-Smith 2000 is also related; it considers how redistributive policy is affected by institutional features. Specifically, it predicts that if legislatures are determined by multiparty proportional representation and engage in legislative bargaining, there will be more redistribution than under a majoritarian two-party system. Feddersen and Gul 2014 propose a model in which (policy-motivated) parties' probability of winning depends on both voter and donor support. A win by the conservative party causes an increase in inequality and a shift of the donor distribution of ideal points to the right, creating persistence in rises to inequality. While they also consider socio-economic conflict and the dynamics of inequality, the focus is on the role of campaign finance as the intertemporal linkage.

[4]In the former, case over they value different types of public goods differently, in the latter they disagree over the level of spending.

1.2 The Two-Period Model

There are two periods, $t \in \{1, 2\}$, and two social groups P and R. In every period they are represented by a party, $J \in \{P, R\}$, which shares their preferences. The size of the population is normalised to 1 and group P makes up a fraction $\alpha > \frac{1}{2}$ of the population.

The incumbent government in period t chooses a proportional tax rate $\tau_t \in [0, 1]$, a level of spending on public (non-storable) consumption $G_t \geq 0$, and a level of public investment $I_t \geq 0$. Public consumption and investment must satisfy a government budget constraint:

$$G_t + I_t \leq \tau_t(\alpha y^P(K_t) + (1 - \alpha)y^R(K_t)) = \tau_t y^m(K_t) \tag{1.1}$$

where $y^m(K_t)$ is the aggregate income.

Public investment goes toward maintaining and increasing the stock of public capital K_{t+1} according to a standard law of motion,

$$K_{t+1} = (1 - d)K_t + I_t \tag{1.2}$$

I allow for disinvestment; however, I require that choices over levels of public capital are non-negative, $K_t \geq 0$.

Public consumption is enjoyed equally by all citizens. Public capital influences the (per capita) income of each group, $y^J(K_t)$, for $J \in \{P, R\}$. The two groups may benefit from public capital to a different extent (or lose income, I don't restrict $y^J(K)$ to being increasing).

Governments are able to alter the wage incomes of the two groups through public investments in things like education, health or public transport (which lead to a more skilled, healthy and mobile workforce). Policies that affect social mobility also lead to this type of dynamic redistribution, by changing the expected future income of individuals who belong to a given socio-economic group today.

I assume the following quasilinear period utility for an agent of type J:

$$u^J(G_t, \tau_t, K_t) = \phi(1 - \tau_t)y^J(K_t) + v(G_t) \tag{1.3}$$

The function $v(\cdot)$ is increasing and concave. ϕ is a positive constant, and it describes the marginal utility of private income. Higher ϕ implies a higher cost of raising government revenues for both groups. Agents discount future utility at rate $\beta \in [0, 1]$.

At the start of each period there is an election, and the outcome of the election is endogenous. Specifically, I assume that parties are not able to commit to platforms, and are therefore expected by voters to choose the party's ideal policy, denoted $(\widehat{\tau}_t^J, \widehat{G}_t^J, \widehat{K}_{t+1}^J)$, once in office. The election is decided by the median voter, who I assume to be of the P group (i.e. I assume $\alpha > 1/2$).

In addition to caring about policies, voters care about some other orthogonal dimension of the candidate's ability, leadership or charisma. I model this as an exogenous random aggregate shock for the candidate from the P party. I assume that the shock has the following distribution:

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

where b is a function of the incumbency state. That is, I assume there is a popularity bias in the election that is systematically related to whether the candidate held office in the last term or not. Formally, let me denote the incumbent in period $t-1$ as $\tilde{J}_{t-1}$, then the bias in the election taking place at the start of period t satisfies

$$b = \widehat{b} I\{\tilde{J}_{t-1} = P\} - \widehat{b} I\{\tilde{J}_{t-1} = R\}$$

where $\widehat{b} \in \mathbb{R}$.

This implies an incumbency-related bias. Suppose $\widehat{b} > 0$, the case of a pro-incumbent bias: if the incumbent is P, then the first indicator function equals 1 and the distribution of P's popularity shock has a positive mean, P is more likely to get a boost to his electoral chances, while when the incumbent is R it is the second indicator function that is 1 and P's popularity shock has a negative mean. When $\widehat{b} < 0$ the bias will systematically favour the challenger.

The assumption of a uniform distribution of the shock is not entirely innocuous. In Appendix A.4, I explore in more detail how results can be affected when different distributions are used.

The median voter will prefer the candidate from the P party if and only if the sum of policy utility (including the expected future utility induced by each party's current policies) and the taste shock is higher for party P than for party R.

I solve for the subgame perfect Nash equilibrium of the game with no commitment to platforms.

1.2.1 Second Period - Redistributive Conflict

In period 2, both parties find it optimal to fully disinvest any remaining stock of public capital because it is the last period: $\widehat{I}_2^J = -(1-d)K_2$ for $J \in \{P, R\}$. Note that this is a feature particular to this model with a termination date, only in the final period will the two parties be in full agreement over investment decisions. In sections 1.4 and 1.5, I explore how results change in a fully dynamic model when parties will face repeated conflict over investment strategies.

The government budget constraint must bind in every period, all revenues raised from taxation are spent and governments can't run deficits: $G_2 - (1-d)K_2 = \tau_2 y^m(K_2)$. Raising government revenue means reducing the private income of both groups, if some government revenues go unspent then whoever is in office could do better for their group by lowering the tax rate so as to only just cover expenditures[5].

Substituting for the tax rate from the binding budget constraint, we have

$$\max_{G_2} \phi \left(1 - \frac{G_2 - (1-d)K_2}{y^m(K_2)}\right) y^J(K_2) + v(G_2) \tag{1.4}$$

subject to a feasibility constraint, $G_2 - (1-d)K_2 \leq y^m(K_2)$ (equivalent to the constraint that the tax rate cannot exceed 1).

The optimal level of public consumption, $\widehat{G}^J$, (at an interior optimum) is characterised by the following first order condition:

$$\phi \frac{y^J(K_2)}{y^m(K_2)} = v'(\widehat{G}_2^J) \tag{1.5}$$

The second order condition is satisfied given concavity of the utility of public consumption.

Each government type chooses the level of G_2 that equates the marginal benefit to its group's marginal cost, which is the share of each pound of revenue raised from its group. The group with higher income will, if in power, implement lower public consumption since they bear a higher burden for financing every unit of G. Further, the parties' choice of public consumption is a function only of the ratio of incomes.

I will denote the ratio of the income of the richer group relative to that of the poorer group by $\iota(K) = \frac{max\{y^P(K), y^R(K)\}}{min\{y^P(K), y^R(K)\}}$; and I will refer to this as the level of inequality. Note that while I have labeled the two groups as P and R, the ranking of incomes may be different at different levels of investment so that, for some K, P

[5] The absence of private savings guarantees this, if individuals were able to save privately, there might be a strategic incentive to tax and burn the revenues in order to affect tomorrow's distribution of gross income.

may in fact be the richer group. I will assume that P are poorer in the absence of public capital, $y^P(0) < y^R(0)$.

Lemma 1 describes the conflict in choices over public consumption that arises in the presence of inequality.

Lemma 1. *The group with lower income will spend more on public consumption, and the difference in choices of public consumption is increasing with income inequality. Formally, whenever $y^R(K_t) > y^P(K_t)$, we have $\widehat{G}^R(K_t) < \widehat{G}^P(K_t)$; further, $\widehat{G}^P(K_t)$ is higher whenever $\iota(K_t)$ is higher, and $\widehat{G}^R(K_t)$ is higher whenever $\iota(K_t)$ is lower. The converse is true when $y^R(K_t) < y^P(K_t)$.*

This result is entirely intuitive. Using proportional taxes to fund pure public goods implies redistribution: while everyone enjoys the goods, the costs are borne more heavily by the rich. If preferences over disposable income are linear then the marginal disutility of funding public expenditures depends only on the price faced and not on the income level of the individual. The richer group is therefore less inclined to provide public goods as generously as the poor would desire. Differences in gross incomes translate into a conflict over redistribution, and this conflict is larger the greater the income gap between rich and poor.

In the election at the start of the second period, the median voter will prefer the candidate from the P party if and only if

$$\phi(1-\widehat{\tau}^P(K_2))y^P(K_2)+v(\widehat{G}^P(K_2))+\epsilon_2 \geq \phi(1-\widehat{\tau}^R(K_2))y^P(K_2)+v(\widehat{G}^R(K_2)) \quad (1.6)$$

where $\widehat{\tau}^J(K_2) = \frac{\widehat{G}^J(K_2)-(1-d)K_2}{y^m(K_2)}$.

The probability of P winning for any pair of platforms is

$$p_2 = \frac{1}{2} + \psi\Big\{b+\phi(1-\widehat{\tau}^P(K_2))y^P(K_2) + v(\widehat{G}^P(K_2)) - \phi(1-\widehat{\tau}^R(K_2))y^P(K_2) - v(\widehat{G}^R(K_2))\Big\} \quad (1.7)$$

which one can also write as

$$p_2 = \frac{1}{2} + \psi\left[b + \Delta^P(\widehat{G}^P(K_2), \widehat{G}^R(K_2))\right] \quad (1.8)$$

where Δ^P is the difference in utility for a citizen of type P between the party of his type winning and the R party winning office.

A party who is in office in period 1 will be able to manipulate, among other things, its probability of being re-elected, through the investment choice. The choice of K_2 can alter the level of inequality and the redistributive conflict associated with it. Specifically, the overall probability of the P-type party winning the second election is increasing with the level of inequality at the start of the second period. So in addition to the exogenous electoral bias from incumbency, there is an endogenous source of electoral advantage which, I show, depends exclusively on the level of inequality.

Lemma 2. *The probability p_2 of party P winning the second election is increasing in $\iota(K_2)$ and is constant with respect to aggregate income $y^m(K_2)$.*

Proof in the appendix.

There are two effects which reinforce each other to improve P's electoral advantage when inequality increases. First, R's position on how much G to provide becomes more distant from what the median voter would like (note that this is true whether P is still the poorer group or whether it overtakes R to be the richer group). The second relates to the difference in tax liabilities faced under the two groups. For example, suppose that P remains the poorer group in equilibrium. Then there is one advantage to R taking office for P citizens - R spends less, which means a smaller tax bill. However, with higher inequality, the price of public spending faced by P is lower so that the difference in overall tax liabilities under a P and R government are smaller anyway. Therefore, higher inequality again helps make the P party more attractive to the median voter. If P are the richer type then the P government chooses lower spending, and the difference in tax liabilities increases with inequality, again increasing the appeal of the P party.

In contrast, changes to aggregate income (while holding inequality constant) do not affect the parties' choice of public consumption nor the tax liabilities faced under each government, so the difference in overall utilities is also unchanged.

1.2.2 First Period - Strategic Investment

Consider the problem faced by government of type P:

$$\begin{aligned}
&\max_{\tau_1, G_1, K_2} \phi(1-\tau_1)y^P(K_1) + v(G_1) \\
&+ \beta\Big\{p_2(K_2)\left[\phi(1-\widehat{\tau}^P(K_2))y^P(K_2) + v(\widehat{G}^P(K_2))\right] \\
&+ (1-p_2(K_2))\left[\phi(1-\widehat{\tau}^R(K_2))y^P(K_2) + v(\widehat{G}^R(K_2))\right]\Big\}
\end{aligned} \tag{1.9}$$

subject to the feasibility constraint

$$y^m(K_1) \geq G_1 + K_2 - (1-d)K_1$$

The first order conditions for public consumption, at an interior solution is:

$$\phi \frac{y^P(K_1)}{y^m(K_1)} = v'(G_1) \tag{1.10}$$

Using the notation $y^J_K(K) = \frac{\partial y^J(K)}{\partial K}$, that for public investment is:

$$\begin{aligned}\frac{\phi}{\beta}\frac{y^P(K_1)}{y^m(K_1)} =& \phi y^P_K(K_2) + \phi(1-d)\frac{y^P(K_2)}{y^m(K_2)} \\ &+ \phi\left(p_2\widehat{\tau}^P(K_2) + (1-p_2)\widehat{\tau}^R(K_2)\right) y^P(K_2)\left(\frac{y^m_K(K_2)}{y^m(K_2)} - \frac{y^P_K(K_2)}{y^P(K_2)}\right) \\ &+ \frac{\partial\iota(K_2)}{\partial K}\left[\frac{\partial p_2}{\partial\iota_2}\Delta^P(\widehat{G}^P_2,\widehat{G}^R_2) + (1-p_2)\frac{\partial\widehat{G}^R_2}{\partial\iota_2}\left(-\phi\frac{y^P(K_2)}{y^m(K_2)} + v'(\widehat{G}^R_2)\right)\right]\end{aligned} \tag{1.11}$$

I assume the second order condition is satisfied[6]. The left hand side is again the marginal cost of public expenditures. The right hand side of the expression is the (expected) marginal benefit of public investment.

Choosing investment policies involves considering the following effects:

- Private income - the more a group's own income grows with public capital the more it will invest.
- A price effect - taking as fixed the level of spending tomorrow, changes to public capital affect the price of public funds and hence, the overall tax liability[7]. The

[6]The second order condition for P's choice of next period public capital is:

$$\begin{aligned}\frac{\partial^2 u^P}{\partial K_2^2} =& \beta\Bigg\{\phi\left(1 - p\widehat{\tau}^P(K_2) - (1-p)\widehat{\tau}^R(K_2)\right) y^P_{KK}(K_2) \\ &- \phi\left(p\frac{d\widehat{G}^P(K_2)}{dK} + (1-p)\frac{d\widehat{G}^R(K_2)}{dK}\right)\frac{y^P(K_2)}{y^m(K_2)}\left(\frac{y^P_K(K_2)}{y^P(K_2)} - \frac{y^m_K(K_2)}{y^m(K_2)}\right) \\ &+ \phi\left(p\widehat{\tau}^P(K_2) + (1-p)\widehat{\tau}^R(K_2)\right) y^m_K(K_2)\frac{y^P(K_2)}{y^m(K_2)}\left(\frac{y^P_K(K_2)}{y^P(K_2)} - \frac{y^m_K(K_2)}{y^m(K_2)}\right) \\ &+ \phi\left(1 - d + y^m_K(K_2)(p\widehat{\tau}^P(K_2) + (1-p)\widehat{\tau}^R(K_2))\right)\frac{y^P(K_2)}{y^m(K_2)}\left(\frac{y^P_K(K_2)}{y^P(K_2)} - \frac{y^m_K(K_2)}{y^m(K_2)}\right) \\ &+ \phi\frac{y^P(K_2)}{y^m(K_2)} y^m_{KK}(K_2)(p\widehat{\tau}^P(K_2) + (1-p)\widehat{\tau}^R(K_2)) \\ &(1-p)\frac{d^2\widehat{G}^R(K_2)}{dK_2^2}\left(v'(\widehat{G}^R_2) - \phi\frac{y^P(K_2)}{y^m(K_2)}\right) \\ &(1-p)\frac{d\widehat{G}^R(K_2)}{dK_2}\left[v''(\widehat{G}^R_2)\frac{d\widehat{G}^R(K_2)}{dK_2} - \phi\frac{y^P(K_2)}{y^m(K_2)}\left(\frac{y^P_K(K_2)}{y^P(K_2)} - \frac{y^m_K(K_2)}{y^m(K_2)}\right)\right]\Bigg\} < 0\end{aligned} \tag{1.12}$$

[7]This can be broken down into two effects: public capital affects the private income that is

richer group will invest less if investment raises the price of public funds, i.e. increases inequality, and will invest more if it reduces the price of public funds, i.e. reduces inequality. The converse is true for the poor group.

- Electoral advantage - P will gain from investing if this increases inequality, and redistributive conflict, making R less attractive to the median voter tomorrow. In turn, R will have a higher incentive to invest when it reduces inequality, making its party more appealing to voters tomorrow due to the lower redistributive conflict.

- Strategic effect - since with positive probability a group in power will fail to get re-elected, they will have an interest in reducing the future level of redistributive conflict. This will bring the policy of the other group next period closer to their own.

- Recovered costs - undepreciated public capital can be disinivested tomorrow, partially recovering the cost of the original investment.

My main result is that if the incumbent has a stronger ex-ante electoral advantage, then it will choose investment such that it allows for higher inequality in period 2. I denoted the degree of incumbency advantage by $\widehat{b}$.

Proposition 1. *A party that enjoys a stronger incumbency advantage will invest less in reducing inequality:* $\left(\frac{\partial \widehat{K}_2^J}{\partial \widehat{b}}\right) > 0$ *if and only if* $\frac{\partial \iota(\widehat{K}_2^J)}{\partial K} > 0$

Proof in the Appendix.

Proposition 1 says that when an investment is income-equalising, then a higher incumbency advantage induces governments to invest less. When the investment leads to higher inequality, higher incumbency advantage will lead to larger investment levels. So, in summary, a stronger pro-incumbent bias results in higher inequality in pre-tax incomes.

The key effect driving this is the strategic effect that investment has on future redistributive policy by altering the level of inequality. When a party currently in office is less likely to hold on to power, it has a higher incentive to reduce tomorrow's redistributive conflict. By reducing inequality a government representing the interests of the poor can induce a higher level of public consumption by a government that represents the rich (and vice versa).

taxed and it alters the tax rate needed to finance spending because the tax base may change. If a group's income grows at a faster rate than aggregate income as K increases, then the overall tax liability faced by this group will be higher as K increases, i.e. the price it pays for G will rise

Changes to the incumbency bias also affect the marginal benefit of investment in terms of the price effect, but the direction is always the same as for the strategic incentive. Consider the richer group. A lower probability of surviving in office makes the richer group more determined to reduce inequality in order to mitigate overspending by the opposition group, for example. And since it is more likely that the larger spending programme of the two will be implemented, it wants to reduce the price of public spending it faces; again it does so by reducing inequality.

It is worth noting that the marginal benefit of public capital in terms of higher (equilibrium) probability of winning is independent of the incumbency bias ($\frac{\partial^2 p_2}{\partial \hat{b} \partial \iota_2} = 0$). This relies on the particular distributional assumption for the aggregate shock; I discuss how relaxing this assumption changes results in Appendix A.4.

There are another two comparative static results I can obtain: one is about the effect of office rents, the other concerns how sensitive voters are to the platforms of politicians. I modify the model slightly, by introducing r, an office rent enjoyed by the party who wins office in addition to his policy utility. Voter sensitivity to policies is parametrised by ψ from the distribution already assumed earlier. Higher ψ corresponds to smaller variance of the taste shock and, hence, makes electoral outcomes more sensitive to policies.

Proposition 2. *Higher office rents r and/or higher voter sensitivity ψ induce the P (R) party to invest less (more) in reducing inequality.*

Proof in the Appendix.

Higher office rents mean parties give more priority to winning for the sake of it. In that case the P party will try to exacerbate future redistributive conflict while the R party will try to mitigate it, since Δ^P is increasing with inequality.

Higher responsiveness of voting, ψ, increases the effectiveness of inequality-altering investment as an instrument for boosting electoral prospects. Again, for the P party, this makes it more attractive to try to keep inequality high; for the R party instead this increases incentives to invest in reducing inequality.

Finally, I make an observation about the distributional assumption of the popularity shocks. I have endogenised the probability of winning by proposing a model of probabilistic voting where there is some source of uncertainty, for example it could be driven by voters' learning about some orthogonal dimension of quality which they have incomplete information about. In the preceding analysis I assumed that the shocks had a uniform distribution. The implication of this assumption is that the slope of the CDF is constant and, therefore, that shifts to the distribution caused by the incumbency bias do not alter the return to changes in a party's choice of investment in terms of improving their electoral prospects. In other words, the *electoral* incentive is unaffected by the magnitude of the incumbency bias.

If one generalises the distribution of shocks, then this no longer holds true. In Appendix (A.4) I generalise the distribution of shocks and consider how the comparative static with respect to $\widehat{b}$ could be affected by the electoral incentive for investment. I find that there are cases where this effect could push the comparative static in the opposite direction to the other effects. The most interesting example is a situation where the P group may have a CDF that is concave in the region of equilibrium investment; in this case, at low values of $\widehat{b}$, P can benefit a lot from keeping inequality high in order to increase its electoral appeal over the R party, whereas as $\widehat{b}$ increases, the marginal benefit of keeping inequality high to raise the probability of winning becomes rather small. In that case, P may be more willing to invest in inequality-decreasing policies when he has a stronger incumbency advantage, if the electoral effect is so strong as to overcome the other effects mentioned.

If the electoral effect is small enough, my results are robust and the relationship in Proposition 1 carries through with more general distributions of the shock.

1.3 Outcomes under Delegation to a Utilitarian Social Planner

How do outcomes in the strategic model of democracy compare to the situation in which a central planner, who maximises a utilitarian social welfare function, decides policy? Are there inefficiencies arising from strategic effects in the setting with political turnover?

I assume the central planner gives equal weight to every individual of the population, and therefore gives weight α to the utility of the P-type individuals and weight $(1-\alpha)$ to the utility of the R-types.

In the second period, investment is costly and has no value because of the lagged process of capital growth, so she will not invest and otherwise solves

$$\max_{\tau_2, G_2} \phi(1-\tau_2)\left(\alpha y^P(K_2) + (1-\alpha)y^R(K_2)\right) + v(G_2) \tag{1.13}$$

subject to the feasibility constraint

$$G_2 - (1-d)K_2 \leq y^m(K_2)$$

Since all individuals value public consumption equally, and since the utility of private consumption is linear, this implies that the central planner acts as if she maximises the utility of an individual with the mean income of society, y^m.

At an interior solution, the following first order condition must hold,

$$v'(G_2^{cp}) = \phi \tag{1.14}$$

and is sufficient since concavity of $v(G)$ ensures the second order condition holds.

The optimal level of public consumption for the central planner is independent of the public capital stock, and of the level of inequality, and (even for a longer horizon problem) is constant over time. This is because the price that she cares about, the cost of financing a unit of public expenditure, is simply ϕ. I denote the central planner's choice of public consumption by $\bar{G}^{cp}$.

Further, the central planner chooses a level of public consumption that lies between the levels that would be chosen by partisan governments, as long as there is some inequality in society. Her choice of public consumption will be higher than the ideal of the richer group and lower than the ideal of the poorer group:

$$max\{\widehat{G}_t^P, \widehat{G}_t^R\} > \bar{G}^{cp} > min\{\widehat{G}_t^P, \widehat{G}_t^R\}$$

In the first period, the central planner solves,

$$\max_{\tau_1, G_1, K_2} \phi(1-\tau_1)y^m(K_1) + v(G_1) + \beta\left[\phi\left(1 - \frac{\bar{G}^{cp} - (1-d)K_2}{y^m(K_2)}\right)y^m(K_2) + v(\bar{G}^{cp})\right] \tag{1.15}$$

subject to the feasibility constraint $y^m(K_1) \geq G_1 + K_2 - (1-d)K_1$.

As I mentioned above, the optimal level of public consumption is the same in every period, $\bar{G}^{cp}$. The first order condition for public capital is

$$\phi = \beta y_K^m(K_2) + \beta\phi(1-d) \tag{1.16}$$

The central planner's condition for optimal investment is to equate the marginal social cost of investment (net of any recovered costs from undepreciated public capital in the final period) to the marginal social benefit, which takes on a particularly simple form, it is just the discounted gross gain in aggregate income. All other effects present in the strategic game, now vanish. There is no chance that tomorrow's decision will be made by an agent with different preferences, so there is no role for influencing either the chances of remaining in power (electoral advantage effect) or the decision of tomorrow's incumbent (strategic effect). And now there is no price effect either because the aggregate price doesn't change with public capital.

Appendix A.5 shows the central planner's investment choice in the infinite horizon problem, where again she will be exclusively concerned with the aggregate growth effects of investment. The feature that investment decisions depend only on aggregate income growth carries through.

In a sense, the central planner solves a commitment problem: the two parties are worried about uncertainty over how G will be chosen in future, and use public capital in part as an instrument to mediate this conflict. The only role for public capital in the central planner problem is if it serves to expand the pool of income available to be redistributed, but redistribution itself can be done through G, once the size of the pie is established.

While the ranking of public consumption choices by ideological parties and a utilitarian central planner is unambiguous, whether political agents invest more or less than the central planner is less clear.

Consider the first order condition for the P government

$$\begin{aligned}\frac{\phi}{\beta}\frac{y^P(K_1)}{y^m(K_1)} &= \phi y_K^P(K_2) + \phi(p_2\widehat{G}_2^P + (1-p_2)\widehat{G}_2^R)\frac{y^p(K_2)}{y^m(K_2)}\left(\frac{y_K^m(K_2)}{y^m(K_2)} - \frac{y_K^P(K_2)}{y^P(K_2)}\right) \\ &+ \frac{\partial\iota(K_2)}{\partial K}\left[\frac{\partial p_2}{\partial\iota_2}\Delta^P(\widehat{G}_2^P,\widehat{G}_2^R) + (1-p_2)\frac{\partial\widehat{G}_2^R}{\partial\iota_2}\left(-\phi\frac{y^P(K_2)}{y^m(K_2)} + v'(\widehat{G}_2^R)\right)\right]\end{aligned} \tag{1.17}$$

I inspect the conditions for optimal investment of the P group and the central planner. A P-type government faces a different marginal cost (it could be higher or lower). The marginal gain in gross income may be higher or lower than for P, even if we know the investment's effect on income inequality (this only tells us the rate of growth). If the investment is income equalising, the price effect is negative (investing makes future expenditures more costly for P-types), the electoral advantage effect is also negative (lower inequality makes the median voter closer to indifferent between a P and R candidate) and the strategic effect is positive (lower inequality insures P-types partially against government turnover). Therefore one cannot say overall, whether the P party would invest more or less than the central planner. The same ambiguity arises for the R-type party so that we cannot rank its investment choice either.

One can reach some general conclusions about how the investment plans of the two types of agents compare. The central planner would prescribe high investment whenever it will help raise aggregate output, and low investment whenever there is a weak output growth effect. Parties instead are likely to be concerned about the distributional effects of investment. For investments that are strongly growth-enhancing but increase inequality, political agents will generally under-invest, because of the commitment problem. In contrast, they will tend to over-invest in forms of public capital that, while having more modest aggregate benefits, are effective at reducing inequality. This is reminiscent of the "political failure" described in Besley and Coate 1998.

In appendix A.6, I present an example that illustrates the arguments above: I consider the case of an investment that is purely redistributive toward the P group, with zero aggregate output effect. Clearly, the central planner would not invest at all. However, as long as the threat of turnover is high enough, both parties will choose to invest positive amounts. This investment is inefficient, because it is costly and does not expand the pool of available resources in future. The inefficiency comes from the commitment problem; the two parties cannot make a binding promise of moderate future redistribution, so they try to bring the rival party closer to its own redistributive preferences, and they're willing to pay a cost today for that insurance.

The analysis of appendix A.6 highlights another interesting point. I consider how expected welfare changes with the incumbency bias. Assuming that the identity of the initial incumbent is equally likely to be either party, I find that both groups would be better off in a world with a weaker bias toward the incumbent[8].

Threat of alternation helps to rectify one inefficiency through another. The first inefficiency is that partisan political groups do not correctly internalise the effects of static redistribution: the poor under-weigh the costs to society of public consumption, the rich under-weigh the benefits. Hence, parties' choices of G move away from $\bar{G}^{cp}$, and the bigger the income gap the bigger the departure from socially optimal transfers. Threat of alternation induces parties to try and close the gross income gap, reducing the variance in G and approaching more efficient levels of transfers.

Higher threat of turnover involves a trade-off in welfare, since it will lead to higher expenditures on K but reduce the variance in G. When reducing inequality can be done at low cost, it will be welfare-increasing overall to have a lower probability of re-election.

[8]In Section 1.5, I show that in a fully dynamic environment, groups will be better off with lower incumbency advantage, even conditional on starting out as the incumbent.

1.4 A fully dynamic framework - Markov Perfect Equilibria

In section 1.2, I analysed a two period model which highlighted a strategic incentive for either group to invest in income-equalising forms of public capital, in order to reduce conflict over within-period redistribution in the final period. In that setup, with a termination date to the game, proportional taxation and preferences given by quasilinear utility, I found a clean relationship between the degree of incumbency advantage and the amount of investment, which depended exclusively on whether investment reduced or increased inequality.

Next, I explore to what extent my results extend to a framework with no termination date, and what new effects arise when there are repeated investment decisions. The general framework I would like to analyse is a simple extension of my earlier framework to an an infinite horizon, and I will focus on exogenous election probabilities for simplicity. Let $p \in (0, 1)$ denote the incumbent's probability of being re-elected.

Individuals' preferences are still given by the period payoff

$$u_t^J = \phi(1 - \tau_t)y^J(K_t) + v(G_t)$$

Each group will maximise the discounted sum of period payoffs, $\sum_{t=0}^{t=\infty} \beta^t u_t^J$ by choosing policy whenever they are in office subject to the feasibility constraint,

$$y^m(K_t) \geq G_t + K_{t+1} - (1 - d)K_t$$

The choice of public consumption is a purely static problem. Assuming an interior solution, if party $\tilde{J}$ is in power it will choose according to the usual first order condition:

$$v'(\widehat{G}^{\tilde{J}}(t)) = \phi \frac{y^{\tilde{J}}(K_t)}{y^m(K_t)} \tag{1.18}$$

Let me substitute for the optimal public consumption choice of the incumbent, as well as for the tax rate from the binding budget constraint of the government.

$$u_t^J = \phi \left(1 - \frac{\widehat{G}^{\tilde{J}}(K_t) + K_{t+1} - (1 - d)K_t}{y^m(K_t)} \right) y^J(K_t) + v(\widehat{G}_t) \tag{1.19}$$

As long as the solution to the problem of the incumbent is such that the tax rate is less than one, then it is also the solution to the modified problem in which public consumption choices are assumed to be those satisfying (1.18), and the incumbent only chooses an investment level (i.e. K_{t+1}). I focus on analysing the problem with a single choice variable.

The problem is a stochastic game in the sense of Fudenberg and Tirole 1991; the history at each period is summarised by a state. In my setting, there are two state variables: one is the identity of the incumbent, the other is the level of public capital. Since the probability of winning tomorrow's election is p for the current incumbent and $(1-p)$ for the current opposition party, this state variable has the Markov property. The law of motion of public capital is deterministic and tomorrow's capital depends only today's action and current level of capital, so it is also a Markov process. Finally, current payoffs depend only on the state and the current action, which is investment or equivalently a choice over tomorrow's capital: $u_t^K = u^J(K_{t+1}^{\tilde{J}}; K_t)$

Players' strategies must specify a plan of action at each state at which they are called upon to make an investment decision, and in principle this strategy could be a function of the entire history (perfect equilibrium), $\sigma^J(h_t)$. Call the space of all possible strategies of a player Σ.

I focus on Markov Perfect Equilibria (MPE), which are subgame perfect and in which strategies depend on history only through the state. More precisely, the strategy of a player is the same for any two histories leading to the same current state, and therefore strategies can be said to be "memoryless"[9]. I can write the strategies as functions of the current capital level, with the understanding that it is the incumbent's choice that will be implemented: $K_{t+1} = \sigma^{\tilde{J}}(K_t)$.

For any general pair of strategies (σ^P, σ^R) and initial state $(\tilde{J}_0, K_0)$, there is a corresponding expected discounted sum of period payoffs for each player. Let me call this the value of the game. The value of the game for group P, for a particular strategy profile, will be denoted $V^I(K_0; \sigma^P, \sigma^R)$ when it starts off as incumbent and $V^C(K_0; \sigma^P, \sigma^R)$ when it starts off as challenger. Similarly, for R, I will write $W^I(K_0; \sigma^P, \sigma^R)$ when it starts off as incumbent and $W^C(K_0; \sigma^P, \sigma^R)$ when it starts off as challenger.

A Markov Perfect Equilibrium is a profile of Markov strategies $(\widehat{\sigma}^P(K), \widehat{\sigma}^R(K))$ such that

$$V^I(K_0; \widehat{\sigma}^P(K), \widehat{\sigma}^R(K)) \geq V^I(K_0; \sigma^P, \widehat{\sigma}^R(K)) \text{ for all } \sigma^P \in \Sigma$$

[9]Focusing on MPE precludes me from studying equilibria in which parties may, for example, punish each other for past behaviour. This could be an alternative way of solving the commitment problem described in Section 1.3.

$$W^I(K_0; \widehat{\sigma}^P(K), \widehat{\sigma}^R(K)) \geq W^I(K_0; \widehat{\sigma}^P(K), \sigma^R) \text{ for all } \sigma^R \in \Sigma$$

While general existence results have been obtained for MPE in stochastic games with a finite number of states and actions, existence results for uncountable state spaces are much harder to obtain. In the remainder of this section, I describe some features of MPE, assuming existence. In section 1.5 I adapt the environment to conform with a well known analytical setup, the linear quadratic framework. While this setup still has limitations in terms of obtaining a full formal characterisation, I am able to find parameter configurations for which equilibria exist. Working with these equilibria, I reconfirm the role of the strategic interaction between static and dynamic elements of redistribution, and I obtain some new insights about the more involved dynamic effects.

The game I analyse has a very similar form to the separable sequential games studied by Fudenberg and Tirole 1991. Crucially, the payoffs are of the form:

$$U_t^J = u_t^J(K_t, K_{t+1}) + w_t^J(K_{t+1}, I_{t+1}, I_{t+2}, ...)$$

That is, the expected continuation value of the game depends only on tomorrow's state, given some profile of future actions. The only difference in this environment is that the identity of movers is stochastic. I will need to be careful about this feature of the environment when applying their monotonicity results.

Monotonicity results for the equilibrium strategies will hold when sorting conditions are satisfied. The sorting conditions require the cross-partial of the stage payoff has a constant sign. This means that the effect of the initial state on the marginal return of the action is itself monotonic:

$$\frac{\partial^2 u^J(K_t, K_{t+1})}{\partial K_t \partial K_{t+1}} \geq 0 \tag{CS$_J^+$}$$

$$\frac{\partial^2 u^J(K_t, K_{t+1})}{\partial K_t \partial K_{t+1}} \leq 0 \tag{CS$_J^-$}$$

Specifically, we have:

$$u^J(K_t, K_{t+1}) = \phi \left(1 - \frac{\widehat{G}^J(K_t) + K_{t+1} - (1-d)K_t}{y^m(K_t)} \right) y^J(K_t) + v(\widehat{G}^J(K_t))$$

This implies that the sorting conditions are pinned down precisely by which group has a higher rate of income growth with respect to public capital, since

$$\frac{\partial^2 u^J(K_t, K_{t+1})}{\partial K_t \partial K_{t+1}} = \frac{y^J(K_t)}{y^m(K_t)} \left(\frac{y_K^m(K_t)}{y^m(K_t)} - \frac{y_K^J(K_t)}{y^J(K_t)} \right) \tag{1.20}$$

The following remark describes how sorting conditions are related to a group's growth rate with respect to capital.

Remark. As long as government expenditures are financed through a proportional tax and preferences are quasilinear, opposite sorting conditions hold for the two parties.

- If group P have higher income growth with respect to K than group R, the (CS_P^-) condition holds for P and the (CS_R^+) condition holds for R.
- If group R have higher income growth with respect to K than group P, the (CS_R^-) condition holds for R and the (CS_P^+) condition holds for P.

One can then apply Fudenberg and Tirole's monotonicity result and conclude that, as long as the effect of public capital on the ratio of the two groups' incomes is monotonic, the reaction functions of the two parties must be monotonic. Further, one of the reaction functions has a positive slope while the other must have a negative slope, the sign of the slopes depending exclusively on which group has the higher income growth rate with respect to K.

Let the strategy of party J, it's choice of next period public capital, be denoted by $\widehat{K}^J(K)$.

Lemma 3. ***Monotonicity of the reaction functions.*** *If the ratio $\frac{y^R(K)}{y^P(K)}$ is monotonically decreasing in K, the reaction function of the P party is downward sloping, and that of the R party is upward sloping.*

$$\frac{d\widehat{K}^P(K)}{dK} \leq 0 \textit{ and } \frac{d\widehat{K}^R(K)}{dK} \geq 0$$

The converse is true when the ratio $\frac{y^R(K)}{y^P(K)}$ is monotonically increasing in K.

Proof in the Appendix.

All the action is coming from the change to the price of public funds. The marginal benefit curve of tomorrow's public capital is independent of the current state. If investment reduces the income ratio, this means it raises the price of investment expenditures for the P party, since at the higher K the P group is relatively richer. Therefore, the P party must prefer lower levels of next period public capital when current levels are higher.

1.5 The Linear Quadratic Framework

Linear quadratic problems, in which the stage payoff is quadratic and the transition law is linear, have well known general forms for the policy functions and value functions. I recast the model in such a way that I will be able to use the solution methods of optimal linear quadratic control. These methods allow for an analytical solution only in the form of an implicit solution to a complex matrix equation; however, it is possible to solve the system numerically for particular parameter configurations.

I assume that the utility of public consumption is quadratic and concave:

$$v(G) = \gamma G - \frac{G^2}{2}$$

where $\gamma \in \mathbb{R}_+$. The transition law for public capital is linear

$$y^P(K) = a + gK$$

where $a \geq 0$ is the income of a P-type citizen when there is no public capital, and $g > 0$ is the growth rate of the income of a P-type with respect to public capital. Crucially, I assume that aggregate income is unaffected by public capital, $y^m(K) = y$. While this assumption removes any effects from aggregate growth, it still lets governments influence future inequality through their choice of investments.

The R group now see their income shrink when investments are made: $y^R(K) = a^R + g^R K$, where

$$a^R = \frac{y}{1-\alpha} - \frac{\alpha a}{1-\alpha} \text{ and } g^R = -\frac{\alpha g}{1-\alpha} < 0$$

I assume $a < y$: the P group are poorer than the R group in the absence of public capital. I also assume that the system starts at the state $K_0 = 0$.

In this context perhaps the language of "investment" and "public capital" may not fit the story so well anymore. I am thinking of actions a government can take to generate a more equal sharing of the pie while not particularly affecting the size of the pie. One interpretation is that these are programmes that increase social mobility, individuals born to families of moderate means now have a higher probability of moving up the ladder, so these policies increase the future expected income of those who are currently poor (and vice versa for the rich). There are many ways in which governments can tackle social mobility. For example, governments can introduce non-discrimination employment laws that will not increase overall levels of employment but instead will increase the chances that a given job post goes to an individual from a more deprived background. They can also offer scholarships that allow low-income students to access better education opportunities, often taking the place of a richer student. These are both examples where actions today can lead to changes in future inequality.

Another interpretation is that these are policies that are primarily redistributive, rather than growth oriented, but are locked in for the future. Certain expansions of the welfare state can be hard to reverse, for example the introduction of a state pension or disability allowance. Once these programmes are in place, they tend to be highly persistent, therefore they could be viewed as a type of long term investment. Finally, certain land reforms could also fall into this category of pure redistribution.

The first observation I make is that this framework requires the interaction of the two groups and presence of strategic effects to generate interior solutions.

Lemma 4. *Interior solutions for P's strategy require the existence of both groups. Interior solutions for R's strategy in any MPE require the existence of both groups and threat of alternation in office; without either of these, R will always choose zero public capital.*

Proof in the appendix.

The setup has a lot of linearity, because the marginal growth rates of income with respect to public capital are constant. In the presence of a single group, the marginal cost of public funds is constant and there is no strategic motive for investment. Either the marginal benefit from higher private income exceeds the cost, and this is true for any additional units of investment so that there is no finite solution, or there is a net marginal cost and no investment takes place (this latter case certainly applies to R, whose private income shrinks with investment).

Interior solutions require a trade-off that brings concavity to the objective function. For R it is the strategic effect that brings concavity because as it invests more, the benefits of reduced social conflict eventually decline. For P, there is additionally the price effect once there are two groups. As it invests more it faces an ever higher price of public funds because inequality decreases. This is why for P the threat of alternation may not be required for an interior solution, although the presence of the R group will, in order for the price effect to exist.

To solve the two player model, I proceed according to the usual methods of optimal control. While earlier I defined the value of the game for a player, for every possible strategy profile, let me now denote by $V^I(K)$, $V^C(K)$, $W^I(K)$ and $W^C(K)$ the value of the game *in equilibrium* for each player and at each initial state (i.e.

evaluated at the equilibrium strategies). For economy of notation, I drop the time subscripts and let today's capital level be denoted by K. The choice of capital chosen for the next period is denoted K'. The policy function of player J, his choice of next period capital K', given today's state K, is denoted by $\widehat{K}(K)$.

I begin by assuming a quadratic form for the value functions of each player (both when they are incumbents and challengers).

$$V^I(K) = A_{VI} + B_{VI}K + \frac{C_{VI}}{2}K^2 \tag{1.21}$$

$$V^C(K) = A_{VC} + B_{VC}K + \frac{C_{VC}}{2}K^2 \tag{1.22}$$

are the value functions of the P type when it is the incumbent and when it is the challenger, respectively. And similarly,

$$W^I(K) = A_{WI} + B_{WI}K + \frac{C_{WI}}{2}K^2 \tag{1.23}$$

$$W^C(K) = A_{WC} + B_{WC}K + \frac{C_{WC}}{2}K^2 \tag{1.24}$$

are the value functions of the type R when it is the incumbent and when it is the challenger, respectively.

The policy functions can be found, in terms of the value function coefficients (A_{VI}, B_{VI} etc), by solving the maximisation problem in a given period, taking as given the form of the continuation value. For example, P solves the problem:

$$\begin{aligned}\max_{G,K'} \phi \left(1 - \frac{G + K' - (1-d)K}{y}\right)(a + gK) + v(G) \\ + \beta \left(pV^I(K') + (1-p)V^C(K')\right)\end{aligned} \tag{1.25}$$

where I have substituted for the tax rate directly from the government budget constraint. The choice of K' is constrained to be feasible (i.e. the tax rate cannot exceed 1) and non-negative: $K' \in [0, y + (1-d)K - G^P(K)]$.

Assuming an interior solution, the choice of public consumption is given by the condition

$$\widehat{G}^P(K) = \gamma - \phi\frac{(a + gK)}{y} \tag{1.26}$$

The optimal investment will be similarly given by the first order condition

$$\widehat{K}^P(K) = \frac{\phi(a + gK)}{\beta y(pC_{VI} + (1-p)C_{VC})} - \frac{(pB_{VI} + (1-p)B_{VC})}{(pC_{VI} + (1-p)C_{VC})} \tag{1.27}$$

Solving R's problem gives us R's policy functions.

$$\widehat{G}^R(K) = \gamma - \phi \frac{(a^R + g^R K)}{y} \tag{1.28}$$

$$\widehat{K}^R(K) = \frac{\phi(a^R + g^R K)}{\beta y(pC_{WI} + (1-p)C_{WC})} - \frac{(pB_{WI} + (1-p)B_{WC})}{(pC_{WI} + (1-p)C_{WC})} \tag{1.29}$$

The policy functions for investment are linear in the state, K. If I substitute the policy function $\widehat{K}$ into the continuation value expression, one can see immediately that the objective function, evaluated at the optimum, is indeed quadratic in the state. That is, I have confirmed that the value function is quadratic, since the maximisation over an objective with a quadratic continuation value in tomorrow's state maps back into the space of quadratic functions.

Now I can write the four Bellman equations that must be satisfied (two for each player, in each of his incumbency states). For each of these to be satisfied at all states K, the constant, linear and quadratic coefficients on either side of the equation must be equal to each other (method of undetermined coefficients). This will yield twelve equations for the twelve unknowns that characterise the value function.

Let us use the notation: $A_V = pA_{VI} + (1-p)A_{VC}$ etc for other similar terms.

The Bellman equation for P in the state where he is incumbent, is:

$$\begin{aligned} A_{VI} + B_{VI}K + \frac{C_{VI}}{2}K^2 = \\ \phi\left(1 - \frac{\left(\gamma - \phi\frac{(a+gK)}{y}\right) + \left(\frac{\phi(a+gK)}{\beta y C_V} - \frac{B_V}{C_V}\right) - (1-d)K}{y}\right)(a+gK) \\ + \gamma\left(\gamma - \phi\frac{(a+gK)}{y}\right) - \frac{1}{2}\left(\gamma - \phi\frac{(a+gK)}{y}\right)^2 \\ \beta\left[A_V + B_V\left(\frac{\phi(a+gK)}{\beta y C_V} - \frac{B_V}{C_V}\right) + \frac{C_V}{2}\left(\frac{\phi(a+gK)}{\beta y C_V} - \frac{B_V}{C_V}\right)^2\right] \end{aligned} \tag{1.30}$$

The other Bellman equations can be similarly written down. They can be found in Appendix A.9, along with the system of twelve simultaneous polynomial equations that I obtain by applying the method of undetermined coefficients.

This system of equations, also known as the matrix Riccati equation, is as far as I can go analytically. It provides implicit solutions for the value function coefficients in terms of the fundamental parameters of the problem $(a, g, y, \gamma, \phi, \beta, \alpha, d, p)$, but the relationship is highly non-linear and cannot be solved analytically.

In the remainder of this section, I focus on particular values of the fundamental parameters, in which case I can solve the matrix Riccati equation numerically. First, I describe a benchmark case with interior solutions. As mentioned earlier, the fact that the R group is willing to sustain positive levels of public capital despite the direct erosion of private income to its group and lack of benefits in terms of aggregate growth, points to a strong strategic incentive to moderate inequality in society. The insurance motive I found in the two-period model therefore extends to a framework with no termination date. Next, I explore several comparative static properties, and provide some further insights on the incentives of the two parties which emerge in this environment with more repeated interaction.

Once particular parameter values are selected, one can always solve the system of equations coming from the method of undetermined coefficients. Of course in order for the output from these equations to indeed be a solution to my problem, I must check whether the following (sufficient) conditions hold for any of the numerical solutions.

1. Second order condition: the continuation value of public capital is concave with respect to K,

$$\frac{\partial^2 \beta(pV^I(K) + (1-p)V^C(K))}{\partial K^2} = \beta(pC_{VI} + (1-p)C_{VC}) = \beta C_V < 0$$

2. Feasibility: the tax rate is non-negative and (weakly) less than 1,

$$\frac{\widehat{G}^J(K) + \widehat{K}^J(K) - (1-d)K}{y} \in [0,1]$$

3. Non-negative public capital: $K^J(K) \geq 0$.

The first condition ensures concavity of the objective function. If it holds, then the first order condition I used to express the optimal choice of investment is a sufficient condition for interior solutions. The last two equations ensure the solution is indeed interior.

The solutions are twelve coefficient values for the four value functions of the problem. In the first example I report the values of all the coefficients and demonstrate that all the required conditions are indeed satisfied. Later I will simply report the features of the solution that are of interest.

1.5.1 Features of the Benchmark Case

The benchmark case has the following parameter values: $a = 50$, $g = 50$, $y = 200$, $\gamma = 200$, $p = 0,35$, $d = 0.5$, $\alpha = 0.5$, $\beta = 0.99$, $\phi = 50$.

The full solution is,

$$\begin{array}{llll} V^I: & A_{VI} = 2.085\ 10^6 & B_{VI} = 91.0 & C_{VI} = 169 \\ V^C: & A_{VC} = 2.080\ 10^6 & B_{VC} = 2065 & C_{VC} = -458 \\ W^I: & A_{WI} = 2.062\ 10^6 & B_{WI} = -1042 & C_{WI} = 144 \\ W^C: & A_{WC} = 2.058\ 10^6 & B_{WC} = 940 & C_{WC} = -482 \end{array}$$

The continuation value of both players is increasing and concave with respect to K' for some range of K' starting at $K' = 0$ (since the value functions are quadratic, the continuation value will eventually decline with respect to K'). The marginal value of the first unit of (next period) capital is $B_V(= pB_{VI} + (1-p)B_{VC}) = 1374$ and $B_W(= pB_{WI} + (1-p)B_{WC}) = 246$ for the P and R parties respectively: investment has benefits for both groups. This marginal benefit declines with the size of the investment as seen from the C coefficients of the value function: $C_V(= pC_{VI} + (1-p)C_{VC}) = -238$ and $C_W(= pC_{WI} + (1-p)C_{WC}) = -263$. Therefore the first order condition for K' is a sufficient condition for an interior solution to optimal investment.

Starting from an initial value $K_0 = 0$ implies a high level of initial inequality, and hence a high price of public funds for the R group, and a low one for the P group. One immediate consequence is that as long as inequality is so high, P will choose a high level of public consumption. Specifically in the first period, it will choose $\widehat{G}^P(0) = 187$ which requires taxing away 93.8% of aggregate income. In contrast R would like a much lower level of public consumption, it sets $\widehat{G}^R(0) = 112$ which requires taxing away 56.2% of aggregate income. This is the static redistributive conflict that I have already discussed.

The P group wish to increase the level of public capital; it both helps it be richer (in terms of pre-tax income) and it has a strategic advantage by ensuring that if it loses office tomorrow, the R group will be willing to finance higher levels of public consumption. Its choice of initial investment is $\widehat{K}^P(0) = 5.7$. The R party is also willing to invest in public capital, despite the erosion to its private income, because it is worried about the high levels of public consumption that P might choose if it takes over. R chooses $\widehat{K}^R(0) = 0.60$ in an attempt to reduce redistributive conflict tomorrow.

In this example, the P party will invest to the point of inverting the ranking of incomes, so that P-type citizens will have higher gross income in the next period than R-type citizens. The R group will of course choose to preserve a higher income than the other group. The P and R labels therefore only allude to poor and rich in the sense of gross incomes *in the absence of public capital*; in equilibrium, the P types will be the richer group in some periods.

All in all, the solutions imply affordable public expenditures: $\tau^P(0) = 0.97$ and $\tau^R(0) = 0.57$, and hence all three conditions for existence are satisfied.

The monotonicity properties are confirmed in the numerical example. The slope of the policy functions has the right sign for each group:

$$\frac{\partial \widehat{K}^P(K)}{\partial K} = \frac{g}{\beta y C_V} < 0 \text{ and } \frac{\partial \widehat{K}^R(K)}{\partial K} = \frac{g^R}{\beta y C_W} > 0$$

This is again due to the effect of K on the price of public funds. As K increases, the P group faces a higher price for public spending and therefore, both the choice of public consumption and public investment will decrease. For the R group, the converse is true; since it faces a lower price for public spending as K increases, it will be willing to spend more on both items.

Both groups are better off starting off as incumbents: $A_{VI} > A_{VC}$ and $A_{WI} > A_{WC}$.

For both groups, the policy function is such that, if they are continuously re-elected they will eventually converge to a level of public capital from which they don't wish to depart: $K^P_{ss} = 5.4$ and $K^R_{ss} = 0.63$. I will call this the parties' steady state levels of capital; however they will only persist as long as the identity of the incumbent doesn't change, if the opponent group enters office it will immediately change the level of public capital. In the steady state of the P group the ranking of incomes is reversed; the two steady states are close to being mirror images of each other, but not quite. P faces the additional cost of maintaining relatively high K and so its steady state tax rate is higher.

To summarise, I have provided a numerical example that proves that solutions to the problem exist for at least some subset of parameter values. The solution described involves the P party choosing high levels of investment, in particular when the current state is low, while the R party choose more moderate but still positive investment. This demonstrates that the strategic effect survives in the dynamic environment. Equilibria in which the R group choose positive levels of investment exist only when the strategic effect is sufficiently strong (for low enough p), and when the redistributive conflict is significant (for high enough γ).

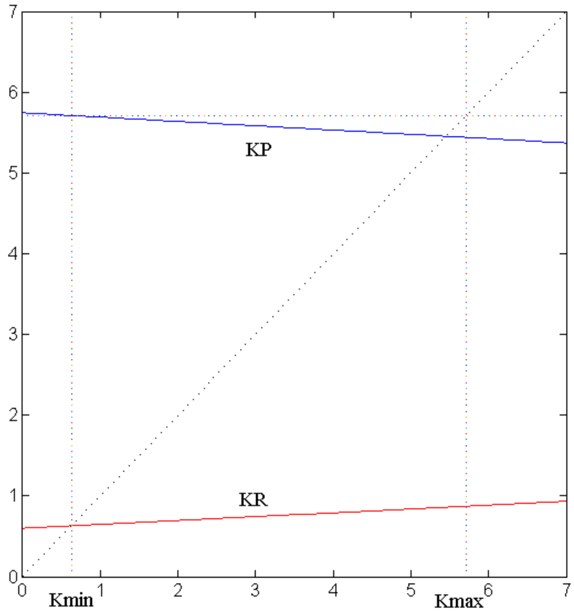

Figure 1.1: The absorbing set, $[K_{min}, K_{max}]$

Absorbing Sets

The benchmark example has one additional feature that applies to all the solutions I work with: the existence of an absorbing set of states. This is a set $[K_{min}, K_{max}]$ such that once $K \in [K_{min}, K_{max}]$, it will remain in the set in all future periods.

As long as both policy functions satisfy, $\frac{\partial \widehat{K}^J}{\partial K} \in (-1, 1)$ then both groups have the feature of converging to a steady state that we discussed above. Further, because of the slopes of the policy functions, P always jumps between choosing K levels above and below K^P_{ss}, while R always approaches K^R_{ss} from whatever side it starts off from. In all the examples I consider $K^P_{ss} > K^R_{ss}$, which implies that the absorbing set is $[K_{min}, K_{max}] = [K^R_{ss}, \widehat{K}^P(K^R_{ss})]$[10]. The absorbing set of the benchmark case is depicted in Figure 1.1.

1.5.2 Comparative statics and changing strategic incentives

I discuss how the equilibrium is affected by the fundamental parameters of the environment, with a particular focus on the effect of the incumbency bias which I identified in Section 1.2 as having a strong relationship with the strategic incentive for investment. In each case, I plot the investment policy functions of the two groups for various values of one parameter, while keeping the other parameters fixed at the values used in the benchmark case.

[10]In principle, it would be possible to have cases where, $K^P_{ss} < K^R_{ss}$ and if so the absorbing state would be $[K_{min}, K_{max}] = [\widehat{K}^P(K^R_{ss}), K^R_{ss}]$

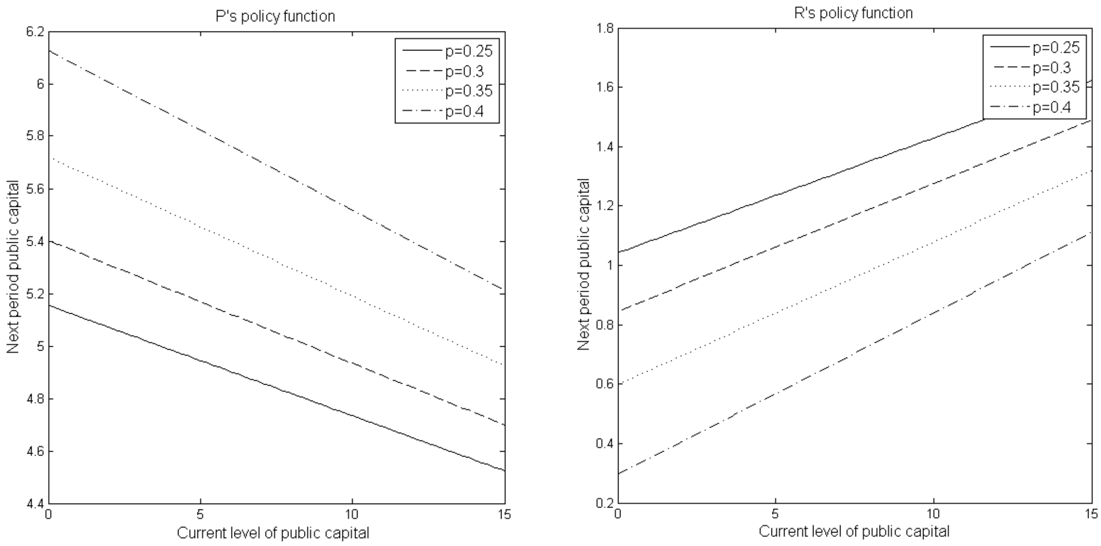

Figure 1.2: Comparative static for the re-election probability, p

The effect of the incumbency bias

I first consider the effect that the probability of re-election of the incumbent has on investment, and show that the result from Proposition 1 still holds in our fully dynamic environment.

Figure 1.2 shows the effect of varying p on the two group's investment policy functions. The choice of investment by the P group increases as it becomes more sure it will stay in office, while the R party invests less when it is more confident it will stay in office.

R is investing because it fears that tomorrow the P party will take over and spend extravagantly if it faces a low price of public expenditure because it is so poor. The more likely the P party is to take over (the lower p is), the more R will invest in reducing inequality, to moderate tomorrow's public expenditures by P were it to take over.

P also tries to moderate inequality more when it is more likely it will lose office. Note that, as I described in the benchmark case, P is choosing high levels of public capital that mean the P-types will have the higher gross incomes in the next period if they get to govern today. Inequality is therefore locally increasing in K' at the equilibrium choice. As p increases the P government is less concerned about losing office, therefore it invests more highly bringing about higher inequality in the next period.

I have identified a set of equilibria for which the comparative static from Proposition 1 carries through, but I cannot claim that this is true for all possible

Change in p	% change in $\widehat{K}^P(0)$ when $d = 0.4$	% change in $\widehat{K}^P(0)$ when $d = 0.5$
0.25 to 0.275	12.2	11.9
0.275 to 0.3	13.7	13.3
0.3 to 0.325	15.4	15.0
0.325 to 0.35	17.5	17.0
0.35 to 0.375	20.0	19.3
0.375 to 0.4	23.1	22.3
0.4 to 0.425	26.9	26.0

Table 1.1: Effect of the depreciation rate on the comparative static for p

equilibria. Now that investment takes place repeatedly today's government will also be concerned about whether it's investments today will be preserved in future. The P government, which aims to maintain higher levels of public capital than R, will be happy to invest more when it is more likely to retain office and not have its costly investment undone in the next period. In equilibria where P overtake R whenever they can choose investment, this effect reinforces the strategic effect. However, in equilibria in which investment is more costly and P does not overtake R, the effects would have opposite signs and it is possible that if the effects are strong enough the comparative static for p is reversed: P may choose to invest more as p increases even though this would reduce inequality.

I can confirm this effect with a simple test. When capital depreciates faster, it matters less that R would undo P's investment, because less capital survives over time in any case. The effect of p on investment driven by the concern for disinvestment reinforces the strategic effect, but it should do so less when d is higher. Looking at the comparative static results when the depreciation rate is 0.4 and 0.5 (see Table 1.1), we do indeed find that the proportional changes in P's choice of investment as p decreases are smaller with the higher depreciation rate[11].

There is one more effect present in the fully dynamic model that was not in the two period model: a party's choice of public capital for tomorrow will also influence tomorrow's investment policies (as well as public consumption policies). P would like R to invest more in general, therefore the more likely it is that R will take over, the more he will want to encourage investment, which he can do by reducing the price R faces (i.e. by investing more). This effect therefore goes in the opposite direction to the dominant strategic effect which is to influence choices of public consumption.

Since the whole policy function for P shifts up with higher p and that for R

[11]It is appropriate to consider the elasticity of investment with respect to p because the lower depreciation rate also results in higher baseline levels of investment. In other words, what I am testing is whether the slope of P's policy function is steeper when d is lower and more capital survives.

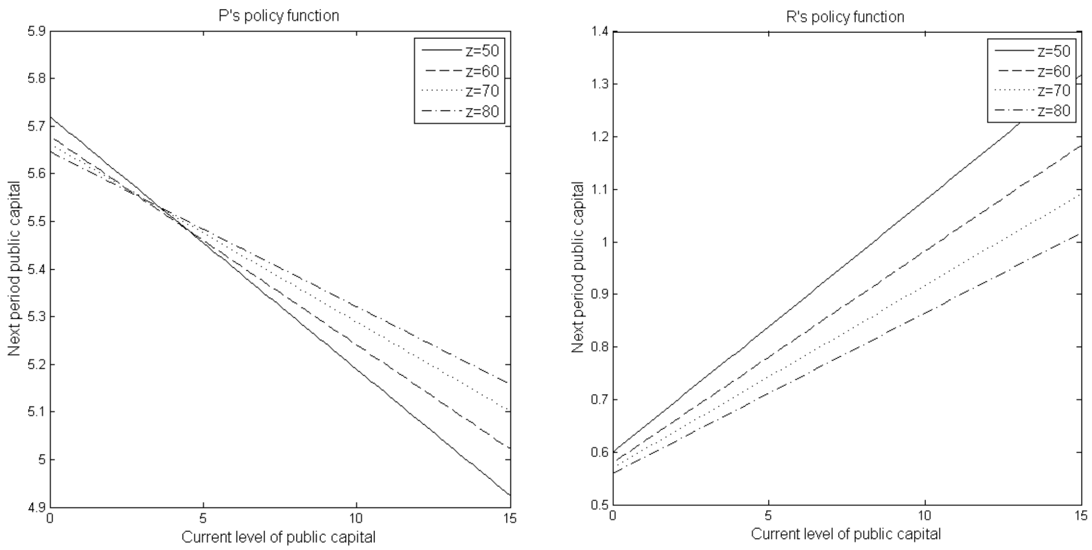

Figure 1.3: Comparative static for the marginal utility of money, ϕ

shifts down, this implies that the absorbing set expands: $K_{min}(= K_{ss}^R)$ is lower with higher p, and $K_{max}(= \widehat{K}^P(K_{ss}^R))$ is higher. A stronger incumbency bias therefore also has the effect of increasing the variance of public capital, and consequently the variance in payoff outcomes for the two groups.

Finally higher p leads to lower A_I coefficients in the value functions of both groups. A is the constant term in the value function and since the economy starts at $K_0 = 0$, it is the lifetime value for each group. Both groups in society end up worse off when there is less turnover in office. This is because when p is higher, the two parties' policy choices are more different from each other. This is the same intuition we discussed in Section 1.3: threat of turnover leads parties to mimic the central planner outcomes to some extent by reducing the socio-economic conflict and hence the variance in period utility.

The effect of the marginal utility of money

I next consider the effect of parameter ϕ on investment incentives. Figure 1.3 shows how the investment policy functions of the two groups are affected as ϕ grows larger. The effect on R's investment choices of higher ϕ is to shift the entire policy function down. Instead for P there is a rotation of the policy function: at low states, P will choose lower next period capital as ϕ increases, but at higher states he will switch towards larger investments.

Higher ϕ makes it more costly to raise government revenues in terms of utility losses from reduced private income. This makes public spending less attractive generally. The main item of expenditure is public consumption; when ϕ grows larger

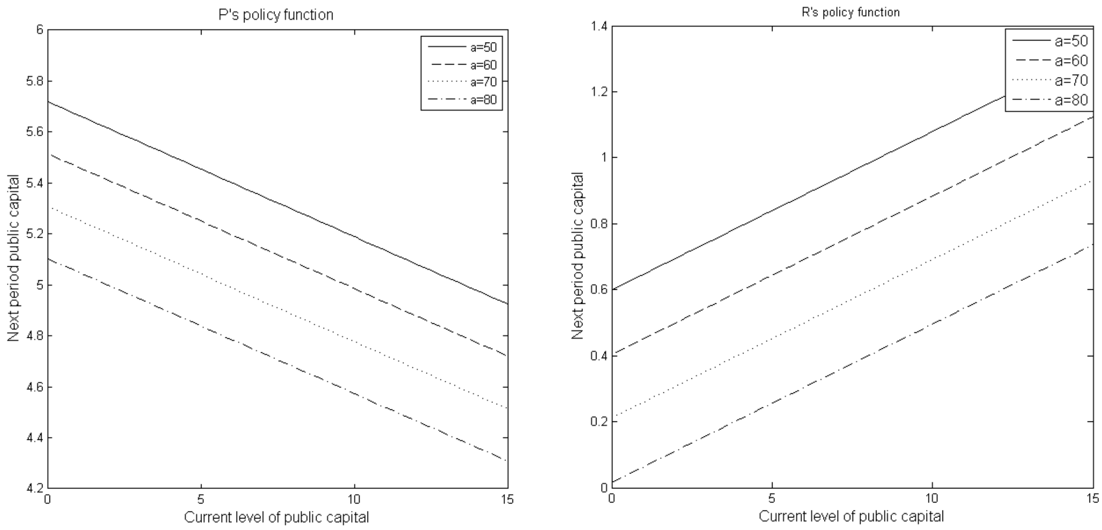

Figure 1.4: Comparative static for the initial income of the poor, a

choices of G shrink and come closer for the two groups, so the static redistribution conflict diminishes. The strategic incentive to invest is weakened and both parties are less worried about keeping inequality low. This is why R invests less, and why P invests more as ϕ increases. However, a higher cost of public funds also makes it unattractive for P to finance the highest levels of investment, this is why its policy function is lower at low initial K where it was investing a lot.

Higher ϕ leads to higher A_I coefficients for both groups. There is a direct effect through making the utility of money higher. Keeping all decisions the same, higher ϕ magnifies the value of income for both groups. There is also a positive effect through G choices, because higher ϕ leads to lower redistributive conflict.

The effect of initial inequality

The level of initial inequality is determined by parameter a, the income of the P group when there is zero stock of public capital. Higher a means that in the absence of any public investment, the two groups have more similar incomes. Figure 1.4 shows the policy functions of the two groups at different values of a. Both groups will invest less as a increases, and the absorbing set will also shift to lower states.

Raising the initial income of the poor has no effect on the slope of the marginal benefit curve of K, it acts purely as a shift of the system, as if it started at a higher initial value of K. Starting off at lower initial inequality and less redistributive conflict means that for both parties the optimum, at which the net marginal benefit of investment falls to zero, is reached with more moderate investment.

Higher a additionally causes both groups to have higher A_I coefficients, so

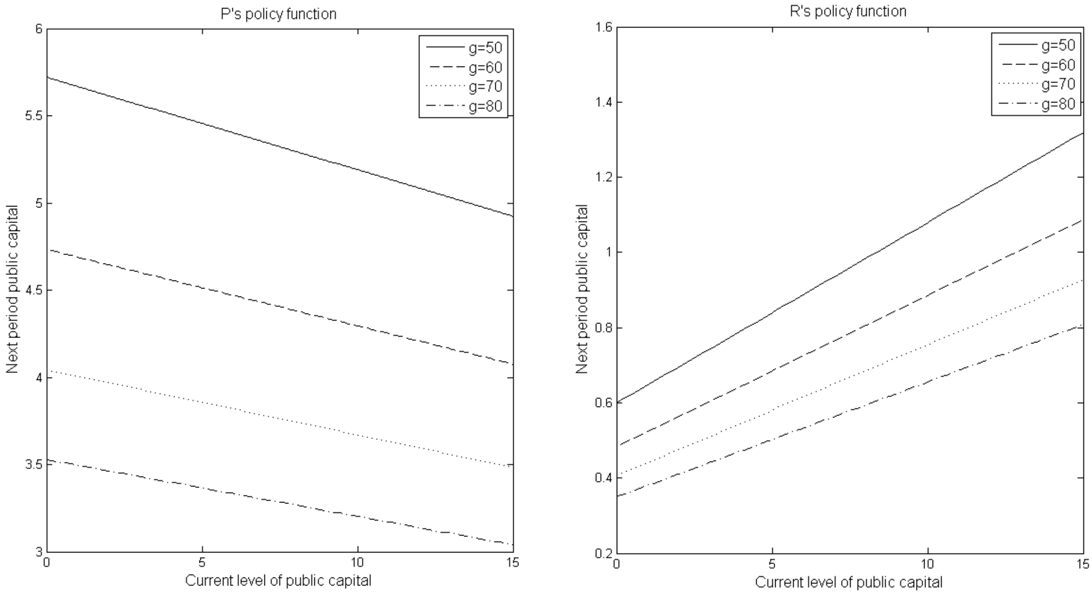

Figure 1.5: Comparative static for the rate of income growth of the poor wrt capital, g

both groups benefit from being in a society that is more equal absent government intervention. This is driven primarily by the reduction in conflict over static redistribution, and further because they are able to achieve low conflict spending less on public capital.

The effect of the growth rate of P's income with respect to public capital

The effect of g, which parametrises how effective public capital is in increasing the P group's income, on investment choices is shown in Figure 1.5. Higher g means that P closes the income gap compared to R (and then pulls away) faster as capital increases. Therefore the marginal benefit of investment also falls faster because redistributive conflict diminishes at a faster rate. Hence both parties, invest less. R can achieve its goal to reduce conflict with lower levels of investment, and P doesn't want to invest too much as it will quickly become too rich for its own good.

Interestingly, while the R-types end up better off as g increases (A_{WI} rises), the P-types end up worse off (A_{VI} falls). The reason is that with higher g the two groups choose lower K levels, and this is true to the extent that although g has gone up, P's average income ends up being lower. This is a slightly counter-intuitive result: when the technology to achieve dynamic redistribution improves, strategic effects could mean that efforts to reduce inequality actually lessen to the extent that the poor end up worse off.

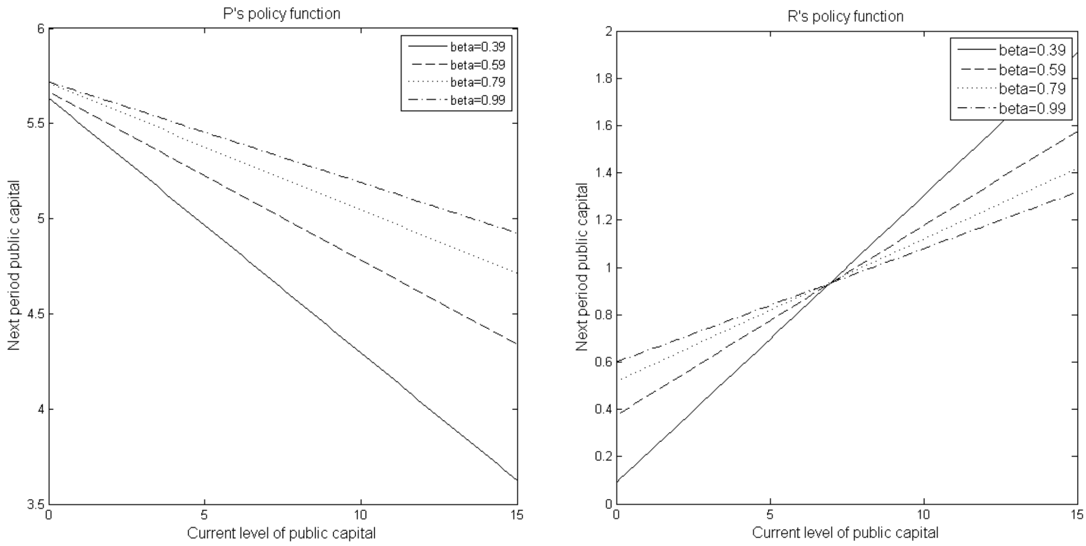

Figure 1.6: Comparative static for the discount rate, β

The effect of the discount factor

A higher discount factor, β, means both groups care more about the future. The effect on investment choices is seen in Figure 1.6: P invests more at any initial K, R invests more at low initial levels but actually invests less starting from high initial K. The main effect is to increase the benefits of investment so that groups choose higher capital in the next period.

For R there is another effect: high K investments lead to lower future income for their group. The marginal benefit curve for them will be higher at low initial K because here the dominant effect is addressing the threat of redistributive conflict for which investment is helpful. But the MB curve will also decline faster, since at higher levels of capital the private income effect comes to dominate. Higher β will therefore mean more investment at low initial K but a flatter policy function because the loss of private income is also given a larger weight when β increases.

The effect of the deprecation rate

A higher depreciation rate d results in lower levels of investment by both groups, as can be seen in Figure 1.7. If capital depreciates faster, this lowers its continuation value. Parties invest because public capital has some value to their group. With higher d however it is more costly to maintain the stock of capital at positive levels period after period, and both parties respond by reducing their investment choices. The absorbing set will shift to lower states.

Higher d makes the P group worse off, while making the R group better off.

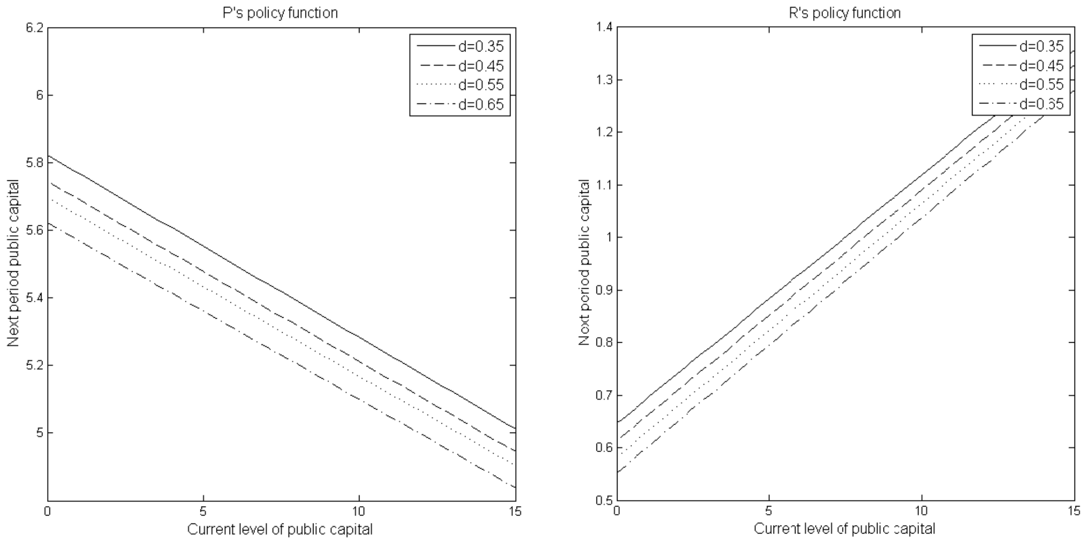

Figure 1.7: Comparative static for the depreciation rate of public capital, d

Since the set of equilibrium levels of public capital is lower when d is higher, P are on average poorer and R are on average richer.

1.6 Conclusions

I have proposed a model in which partisan groups, protecting the interests of the rich and the poor, alternate in political office. The government can redistribute in any period by taxing and spending on public consumption, and it can also make costly investments in altering the future distribution of gross income. Such investments will have a strategic role: under threat of losing office to the opposition, who will not redistribute as desired, incumbents will try to push the opposition's fiscal preferences closer to their own, i.e. to reduce the redistributive conflict in society. The main insight of this work is precisely that the possibility of alternation will drive both parties to engage in dynamic redistribution that will reduce inequality in gross incomes. Efforts to keep gross inequality low will increase with the probability of losing power.

The strategic incentive to keep gross inequality levels low is able to justify support from the party of the rich for programmes that may even do direct harm to its own group, like investments that are purely redistributive and hence, reduce the income of their group. Acemoglu and Robinson 2001 find a similar strategic incentive but in a more extreme setting: political elites may choose to reduce political inequality by extending the franchise, if the threat of revolution is rather high. In both cases, privileged groups willingly relinquish some of their advantage in order to reduce social conflict, but only when some sort of threat of expropriation exists.

My result suggests that political regimes in which a party has a firm grip on power and there is little threat of turnover, will engage less in policies that lead to equalising of incomes, for example investments in human capital. These regimes may undertake large transfer programmes when they wish to protect the poor, but they will rely more heavily on welfare relief and less on building up forms of public capital that help the poor raise their productivity. The analysis of endogenous elections suggests one exception to the main result: large redistributive conflict can help give the party of the poor an electoral advantage, if the returns to this party from raising inequality to increase its electoral lead fall with the size of its advantage, then it may be more willing to reduce inequality when it is more secure in office.

Strategic interactions lead to distortions from the socially optimal investment plan (that which maximises the utilitarian social welfare function), a form of "political failure". This is essentially due to the commitment problem that arises because of uncertainty over who will set policy in future, and is reminiscent of other results obtained in the literature (Alesina and Tabellini 1990, Besley and Coate 1998). For investments that are mostly redistributive with little effect on aggregate income, the results suggest there will typically be over-investment by both parties as they both try to insure against future redistributive conflict. In contrast, there could be under-investment in forms of public capital that significantly increase aggregate output but at the cost of higher inequality. In the absence of a benevolent (utilitarian) dictator however, welfare may increase for both groups when incumbency advantage is lower, by encouraging policies that reduce socio-economic conflict.

Chapter 2

Risk Regulation with a Boundedly Rational Public

Policy makers in charge of regulating risks that are not well understood by the public often face pressure to make decisions which may not be in the public interest, in order to salvage their reputation. I propose a formal model of accountability in which the public update beliefs about risk by a heuristic approach, and retrospectively evaluate the past decisions of policy makers according to current beliefs about risk. Policy makers with good information about the true risks, but concerned with maximising their reputation, will distort their actions away from the efficient policy. This is driven by three different incentives: (i) they will suppress learning in order to prevent beliefs taking a bad turn, creating a bias towards the safe action (tight regulation) (ii) they will pander toward current public beliefs, and (iii) they will try to change beliefs by deregulating if current beliefs don't favour them. The essential trade-off for the policy maker will be between manipulating the beliefs of the public, and aligning his record of policy with beliefs. I consider how some institutional features influence these distortions.

2.1 Introduction

A sign at Piers Park in East Boston reads: "Bicycle Riding, Roller Skating, Skate Boarding, Ball Playing, Kites, Scooters, Wheeled Toys, Swimming, Fishing, Fires, Barbeques, Alcoholic Beverages, Glass Containers, Animals and Unlicensed Vendors are Prohibited." The dilemma faced by regulators who are accountable to the public is illustrated by this seemingly banal example. When perceived risk is low and the authorities regulate, they are condemned by public opinion for being overly cautious. But when they deregulate they allow for the possibility that an accident may happen and they will be blamed for their lax safety standards. A ban on these activities prevents accidents from taking place but also prevents the public from revising their opinions about risks, perpetuating discontent. The regulators face a catch-22 due to the public's wrong priors on risk, how their perceptions of risk are formed and updated in response to events, and the retrospective assessment they face at the end of the day.

Perhaps a more substantial economic example of a regulatory regime that is influenced by public perceptions of risk is provided by the regulation of genetically modified crops (GMOs). In Europe, regulation of GMOs has been much stricter than in the USA, with the EU failing to approve the marketing of many GMOs considered safe in the USA. Since the information available to regulators on both sides of the Atlantic is essentially the same, this is a bit of a paradox. One possible explanation for the different policies lies in the incentives of regulators to please the public. Indeed, public perceptions of the risks of GMOs vary markedly across the two regions: a 1995 survey of consumers reported that only 21% of American consumers regarded genetic engineering as a "serious health hazard"; by contrast, the comparable figure was 85% in Sweden, 57% in Germany, 39% in the United Kingdom, and 38% in France[1].

Pandering to public misperceptions is one possible consequence of exposing regulators to public accountability. I attempt to formalise the mechanism by which individuals form risk perceptions and apply them in their judgement of regulators. The model I analyse provides a rich picture of the incentives of regulators, and of the dynamics of regulatory policy making and public learning. The aim of this paper is to study the interaction of public perceptions of risk and the public accountability of regulators.

We know from the vast literature on risk perceptions, that individuals are liable to make systematic and pervasive mistakes in their assessments of risk. Psychologists have thoroughly documented a range of biases and heuristics applied by individuals

[1] http://ec.europa.eu/public_opinion/archives/ebs/ebs_341_winds_en.pdf.

estimating probabilities for risky prospects. These biases are especially typical in environments where risks are not well understood, for activities that are not part of the day to day experience of the public[2].

Further difficulties arise when the activities involve small probabilities of catastrophic events; learning will be slow and quite likely influenced by the salience of events. As one public health expert puts it: "knowing and understanding the probability of the bad effect and, even more, "appreciating" it can be a challenge... people have problems understanding small probabilities and often set a near-zero probability to zero, thus ignoring it. Or sometimes they fixate on the adverse consequence and ignore the fact that it almost surely will not happen" (Pauly 2007). Both the priors and the updating procedure for beliefs about risk appear to suffer from biases.

This paper contributes to the discussion of risks by considering the interaction between the perceptual model of risk and the institutional environment. When public officials and policy makers are accountable to the public, how do risk perceptions affect the reputation of decision makers and their incentives to choose policy? The accountability framework will bind the hands of public officials and determine what they can and will do to manage risk and perceptions of risk.

This paper also contributes to the literature on inefficiencies caused by systems of accountability. So far the literature has focused on rational models. Maskin and Tirole 2004 have shown that public officials accountable to the (Bayesian) public and whose motivations are uncertain, will tend to pander to popular policies. In Canes-Wrone, Herron, and Shotts 2001, public officials differ in their information about the appropriate policy and may pander, but may also take excessive risks to forge a reputation of being in the know. Dewan and Hortala-Vallve 2013 consider the setting where politician quality determines the success of risky reforms. This means experimenting with reform has learning value, but this value is different for voters and the incumbent, so that he may take too much or too little risk. I depart from these papers by assuming the public are not fully rational, and by considering a fully dynamic environment, which allows me to describe the evolution of beliefs about risk, and of reputations, over time.

There is a question mark over whether the rational model is the most helpful framework to describe and analyse certain policy environments. First, one may be concerned about the cognitive demands involved in solving the public's problem rationally. This would suppose that the public have some underlying prior beliefs

[2]For example, surveys asking individuals to estimate mortality rates for different causes find that for those activities that are considered uncontrollable or evoke feelings of "dread" individuals tend to significantly overestimate risks (Slovic 1987).

about the likelihood of encountering policy makers of different qualities, and also over the possible levels of risk, that they understand how the different types of policy maker will behave in equilibrium for every history of past actions and events and for every state of the world, and that they can correctly update beliefs as new signals are observed. The environment faced by the public in reality is likely to be even more complex: policy makers may differ in their motives as well as their quality, the environment may change over time, there may be other factors that influence risk levels which are not under the control of the policy maker.

There is suggestive empirical evidence, from experiments and from the field, that individuals assessing the performance of politicians and public figures often rely on simple rules. This has been termed low information rationality by Popkin 1994 and judgement heuristics by Tversky and Kahneman 1974: when information is costly to acquire and process, voters' evaluation of politicians will be based on various shortcuts. I discuss this literature further in section 2.2.

My analysis is based on the premise that the public will tend to simplify the procedure by which they evaluate policy makers in such environments. The question a rational individual would want to address is: how good should we expect future performance of this agent to be? This is what will determine the public's decision to reward or retain the decision maker. When the environment is extremely complex, the public may well fail to have a complete model for the information, constraints and motivations of the decision maker, and instead is likely to address a simpler question: how well does he appear to have performed in the past, or how good do his past decisions appear today? Kahneman describes this process as *substitution*.

The updating rule for perceived risk, coupled with the retrospective assessment rule, distorts policy makers' decisions in several ways. On the one hand there is a very strong bias towards strong regulation: even if the policy maker is well informed about the risk and this is minimal, he may find it preferable to prohibit the activity, even in some cases where the public would prefer to allow it. Tight regulation (and in the extreme, banning the activity entirely) constitutes a "safe action" because the decision maker can suppress learning and reduce his reputational risk. In contrast, deregulation always opens up the possibility of an undesirable outcome. Second, there is an incentive to pander to the beliefs of the public. If current beliefs about risk are preserved with high probability regardless of the action taken, then the decision maker will want to conform to those beliefs, even if he knows beliefs to be misguided. Finally, there is an incentive to actively try to change beliefs. This happens when the decision maker's policy record is at odds with current beliefs and is difficult to revise, but beliefs can be aligned by allowing the public to observe events.

The distortions are pervasive, even decision makers with perfect information about the actual level of risk will engage in behaviour to enhance or protect their reputation, causing a loss of public welfare. In many cases, policy gets stuck on a suboptimal course of action. By assumption, my cognitive model implies that learning is not necessarily converging to the truth. Even if the public are able to observe infinite realisations of events, beliefs may remain unstable, with infrequent disasters bringing periods of fear that may be gradually calmed by further exposure to the activity or become entrenched if the decision maker yields to demands for prohibition.

There is no simple remedy for these distortions to policy. In the limits of very low and high levels of risk, increasing the duration of office terms or learning from an "outside economy" will lead to better outcomes. But for intermediate levels of risk these measures can also lead to inferior policy outcomes. Making it costly to switch policy is helpful for preventing over-regulation in response to disasters because regulators cannot change course. But it also increases conservatism ex ante, when the public would support deregulation.

It is worth thinking about the extent to which it is the bounded rationality that is driving results, and to what extent it is the usual agency problem. I address the comparison of the two paradigms, rational and boundedly rational in an extension. However the primary focus of this work is to consider the boundedly rational world, this is the novel contribution to a much explored question of problems in accountability. Just as in any rational model, different assumptions could have been made. My assumptions are not arbitrary, they are motivated by well documented psychological attributes and behaviours. The model produces a rich picture of the incentives of regulators, in particular of the dynamics, which in some cases the rational model struggles to explain.

In section 2.2, I provide a thorough discussion of the literature on both risk perceptions and evaluation shortcuts, and use it to anchor the formal perceptual model. I analyse the effects of the boundedly rational model of accountability on incentives in section 2.4. I go on to consider some variations on the assumptions about the institutional environment in section 2.5. Section 2.6 addresses how results would be different under full rationality, and section 2.7 concludes.

2.2 A Psychological Model of Risk Perception and Performance Evaluation

2.2.1 The literature on cognitive biases

There is substantial evidence that individuals suffer from pervasive biases and rely on simple heuristics when assessing risks, in particular in environments that are unfamiliar and when dealing with small probabilities. Furthermore, there is suggestive evidence that individuals take shortcuts and use simple rules for evaluating the performance of individuals. I address each of these strands of literature in turn.

Perceptions of risk

There is a vast literature in the field of psychology and decision making that tries to understand how people form probability assessments when they have limited information about a stochastic process. Starting with Kahneman and Tversky's seminal work on decision under uncertainty (Tversky and Kahneman 1974), and supported by later studies that built on their work, two departures from rationality have been identified: biases (where beliefs systematically depart from objective probabilities due to the characteristics of the prospect or other irrelevant factors) and heuristics (where individuals do not use available information according to optimal statistical laws to update their beliefs).

Individuals tend to form quick intuitive judgements about difficult questions, such as the risk level of a complex activity that they know little about. When asked to estimate the probability of an event, they often revert to the availability heuristic (Kahneman and Tversky 1973): they will guide their judgement by how easy it is to recollect instances in their memory of the event taking place. This process will be sensitive to salience effects[3]: more vivid and dramatic events may cause larger revisions of probability estimates than less salient ones, and more recent events may carry more weight than those further in the past. This is true even when recency or salience has no bearing on the likelihood of an event today. Studies have also found that other non-relevant elements may influence assessments of risk: a lack of controllability or a sense of dread will often increase risk estimates (Slovic, Fischhoff,

[3]For example, in one experiment K&T ask individuals whether there are more words starting with the letter K or with the letter K in third position. Individuals wrongly answered that those starting with K are more common, because these are easier to retrieve, i.e. more salient in our memory.

and Lichtenstein 1979)[4].

Individuals often favour the content of a piece of information over its reliability or relevance[5]. This preference induces an outcome-bias. After disaster occurs the sense that they have seen disaster with their own eyes, or simple alarmist messages found in the media, are likely to drown out the voices of experts with their technical lingo and statistics. Further, psychologists have recognised a bias towards negative information and experience (Baumeister et al. 2001): bad feedback makes a larger imprint on our impressions. News about disasters will knock our beliefs dramatically and that fear will not be easily erased by facts and figures provided by the authorities; "disagreements about risk should not be expected to evaporate in the presence of evidence" (Slovic 1987).

Another bias reinforces the outcome-bias: a belief in the law of small numbers (Tversky and Kahneman 1973), the confidence individuals have that the frequency of an event over a small number of trials might be a reliable indicator of the underlying probability. Individuals may fail to recognise that extreme frequencies are observed too often in small samples. As a result, after observing a small sample in which no disaster is observed, individuals may infer that the risk is essentially zero; samples with a high a rate of disaster (in the case of small probabilities, a single instance of disaster) may also generate estimates that are too high.

These biases and heuristics are particularly problematic when dealing with unfamiliar environments, which are highly technical and involve small probabilities. As Slovic 1987 puts it, "the mechanisms underlying these complex technologies are unfamiliar and incomprehensible to most citizens. Their most harmful consequences are rare and often delayed, hence... not well suited to management by trial and error learning".

Evaluation of decision makers

There is some evidence from the empirical political economy literature that voters are not fully rational when it comes to voting. For example, individuals' voting decisions are often influenced by information that is irrelevant when it comes to judging either the preferences or skills of candidates. Achen and Bartels 2012 and Healy, Malhotra,

[4]It is important to distinguish between risk perceptions, and judgements about the acceptability of risks. Both of these may be influenced by such factors but I am only concerned with the former in my analysis.

[5]Base rate neglect is a typical example of this; K&T were first to introduce base rate neglect into the literature: individuals overweigh superficial information that feels more representative over statistical information about the population. See for example their experiment in Tversky and Kahneman 1974 where, given a description of an individual, participants were asked to assess the probability of the individual having different occupations.

and Mo 2010 find evidence that incumbents' probability of holding on to office can depend on adverse events that could not have been foreseen or prevented, and for which there was little to be done ex post, such as shark attacks, extreme weather or results in local college football games. They describe the procedure by which politicians are assessed as "blind retrospection": voters are happy with incumbents if they feel things are going well, without deeper understanding of the root causes of outcomes.

Preventing shark attacks may be a tall order on politicians, but improving economic performance is certainly something we expect our politicians to play a role in. However, evidence suggests that voters often rely on naive rules of thumb based on weakly informative criteria. Evidence exists of retrospective economic voting where voters re-elect incumbents during good economic times and oust incumbents in a downturn (Kramer 1971; Arcelus and Meltzer 1975; Bloom and Price 1975); and of the micro counterpart of this, pocketbook voting, where voting is based on an individual's economic circumstances at the time of an election (M. P. Fiorina 1978). Electoral prospects also seem to turn on the most recent conditions rather than the record of success of a politician, and politicians duly respond to this (e.g. Tufte 1980 finds evidence that American governments engage more often in economy-expanding policies close to election time).

The rational basis for economic or pocketbook voting is dubious: the connection between economic shocks and incumbent quality is weak, especially when it comes to fitful events, such as a sudden rise in oil prices or a tumble in the stock exchange. The accountability of politicians is then best described as "rough justice" (M. P. Fiorina 1981): when things seem to be going well voters will maintain confidence in current leaders, and when something doesn't seem right, heads will roll.

Popkin 1994 described voters as being subject to low information rationality: lacking relevant information, or possibly overwhelmed by an excess of information which they struggle to interpret, voters resort to shortcuts when they evaluate politicians. One of the shortcuts he describes is the representativeness heuristic: individuals compare the actions of the politician to their image of what a good politician should be doing.

Finally, as M. P. Fiorina 1978 argues, retrospective voting need not be regarded as "silly": in fact it may be "individually reasonable and systematically desirable". The voters are not following random procedures and are not systematically wrong, they are using simpler rules that are approximately right. Simple rules reduce cognitive costs, provide a clear way for voters to communicate desire for change, and create some good incentives for politicians who will attempt to prevent the negative shocks which are anticipated to cost them their job.

2.2.2 Anchoring the Psychological Model

I propose a cognitive model for how individuals update their perceptions of risk and their judgements about those making regulatory decisions on their behalf. The cognitive model is boundedly rational. Individuals do not have a complete model in mind about the environment they face, specifically regarding the information or objectives that the decision maker may have. Instead, they use simple rules of thumb.

1. *Perceptions of risk.* Individuals' perceptions of risk are modelled on the availability heuristic and are outcome-biased. When assessing how likely a bad event is, they look back over their experience of the activity. If instances of disaster are easy to recall, they will form a high belief about risk; if they struggle to conjure up examples of disaster, their belief will be low. Updating will be adaptive in the face of new observations, and give a higher weight to more recent observations. In the absence of new observations, beliefs will be preserved. The model therefore has a flavour of much of what is discussed in the risk perception literature:

    - A strong outcome-bias;
    - Limited memory, events further in the past carry lower weight despite being equally informative;
    - Underestimation of probabilities when no disasters are observed;
    - Overestimation of events when a disaster is observed;
    - Excess volatility of beliefs over time compared to rational learning;
    - Arbitrary priors, initial beliefs can take any value and may differ substantially from objective probabilities or experts' assessments since they may be driven psychological factors such as controllability or dread.

2. *Evaluation of decision makers.* Individuals engage in a retrospective assessment of past decisions. They do not conjecture the reasons why a decision maker could have made mistakes in the past, to what extent these were driven by greed and special interests, by ineptness of the decision makers or by bad luck. They apply a simple rule of thumb: do past policies chosen by the decision maker match what I think, at this point in time, he should have been doing? I propose a reputation function that penalises deviations from what is today perceived to be good policy[6] and has the following features:

[6]Individuals are assumed to use today's beliefs to judge yesterday's actions, they are not

- A record with more frequent and/or recent[7] deviations from good policy is more heavily penalised;
- Deviations that generate a larger social cost are more heavily penalised (for example, deregulation is more heavily penalised the larger the loss incurred in a disaster and the higher the perceived risk of disaster).

The model depicts an individual who is clearly naive, and in more than one way. Some assumptions may seem more plausible and others more extreme. That individuals may have the wrong prior on risk is strongly supported by many studies of risk perceptions[8].

Perhaps more worrying is the assumption that individuals neglect the strategic information about risk from agents' choices: an activity that is allowed indicates some knowledge on behalf of decision makers. It is not difficult to come up with examples of individuals stubbornly disagreeing with the authorities about the risks they are exposed to: resistance to GM crops, unresponsiveness to measures to improve road safety, deaf ears to warnings about environmental risks. A lack of information about the specific forces influencing public officials' decisions, and more generally low trust in the public process, may help explain why individuals rely heavily on what they see and experience, as more tangible evidence of the risks they face.

Finally, the case for a rough rule of retrospection is based on two observations. First, for many environments it is not feasible for individuals to have a good understanding of the full model with all its complexities. Second, I refer to the literature on retrospection in elections, which strongly suggests that individuals punish those in charge according to rather crude measures of performance. This is supportive of the claim that in complex policy environments the rational model may in fact perform quite poorly in describing people's behaviour and beliefs.

forgiving for the fact that the regulator himself may have also learnt over time. Kahneman 2011 describes how "Once you adopt a new view of the world you lose much of your ability to recall what you used to believe before your mind changed". This effect has been widely documented in experiments where people were asked to express their opinion on a topic, say the death penalty, then hear a set of arguments, and report their views again. When asked to retrieve their previous attitude people are unable to. What is now believed seems, to the individual, to have always been apparent. Fischhoff and Beyth 1975 proved this hindsight bias in a real life example: they asked individuals to estimate the probabilities of several possible outcomes of a visit of Richard Nixon to China and Russia in 1972. He then asked individuals to recall their estimates after the events took place; people systematically overestimated their original guesses for events that were realised and underestimated them for events that were not.

[7] There is substantive evidence that voters focus too heavily on recent rather than cumulative signals of performance; see for example Kramer 1971, Achen and Bartels 2004, or Huber, Hill, and Lenz 2012.

[8] See for example, Slovic, Fischhoff, and Lichtenstein 1979 where beliefs about mortality rates for different causes showed participants' estimates were widly off the mark.

It may seem unnecessary to deal with both cognitive assumptions simultaneously; in fact it would not make sense to use them individually. If voters do not have a complete model in their mind of what is motivating the decision maker, they cannot use the strategic information from policies to deduce anything about either risks or the quality of decision makers. Estimating risks ignoring the strategic information from observed actions goes hand in hand with a non-Bayesian model for quality assessment.

2.3 A Baseline Setup

An activity is to be regulated by a decision maker (DM hereafter), the outcomes from this activity will affect the wider public. The activity generates some gain, a, in every period in which it is allowed to take place; these gains are deterministic and commonly known by all individuals. In each period, with some probability ε the activity results in a disastrous event which causes a large social loss, L; this loss is again well known by all individuals, the probability of disaster however, is not. I assume that all individuals are risk neutral[9].

The DM must choose whether to ban or allow the activity, $y_t \in \{0, 1\}$. When the activity is banned the gain cannot be accrued but the risk is avoided altogether. When the activity is allowed the gain can be enjoyed and an event will take place, $z_t \in \{0, 1\}$ which describes whether or not a disaster occurs (let $z_t = 1$ correspond to the case of disaster). This implies that the welfare maximising strategy is to allow the activity so long as the risk is not too high, specifically so long as $\varepsilon \leq \frac{a}{L} = \hat{\varepsilon}$.

I formalise the cognitive model described in section 2.2. The public's beliefs follow an adaptive law of motion with respect to observed events; when none are observed, they preserve their old beliefs.

$$\tilde{\varepsilon}_t = \begin{cases} \delta\tilde{\varepsilon}_{t-1} + (1-\delta)\, z_t & \text{if } y_t = 1 \\ \tilde{\varepsilon}_{t-1} & \text{if } y_t = 0 \end{cases} \tag{2.1}$$

where $\delta \in [0, 1]$. This implies that the public do not infer anything from the DM's actions about the risk itself.

In addition, at every point in time, the public summarise the DM's policy record through an aggregate measure of his past decisions, $\tilde{y}_t$. The updating rule for the

[9]This is in order to focus on the preferences over uncertainty created by learning and accountability, rather than by explicit tastes for risk.

policy record is again adaptive:

$$\tilde{y}_t = \rho \tilde{y}_{t-1} + (1-\rho)\, y_t \tag{2.2}$$

where $\rho \in [0,1]$.

The public will use a rule of thumb to evaluate the DM, which is based on the degree of alignment between his record of policy and the policy perceived to be correct at that point in time. The reputation payoff at time t is:

$$\Delta(\tilde{\varepsilon}_t, \tilde{y}_t) = (\tilde{y}_t - \mathbb{1}\{\tilde{\varepsilon}_t \leq \hat{\varepsilon}\}) \cdot (a - \tilde{\varepsilon}_t L) \tag{2.3}$$

This reputation function will drive all our results and reflects the retrospective logic described in section 2.2.

To interpret this expression, suppose first that perceived risk is lower than the critical value $\tilde{\varepsilon}_t < \hat{\varepsilon}$. The indicator function takes on the value of 1: the optimal record would be to have always allowed. If the record is $\tilde{y}_t = 1$, then the DM is considered to have made no mistake and he suffers no penalty. Instead if $\tilde{y}_t < 1$, his past actions are considered to be mistaken, and the lower $\tilde{y}_t$ the bigger the mistake. Also, the higher the expected gain from allowing $(a - \tilde{\varepsilon}_t L)$ is, the more heavily they penalise this mistake because the DM is perceived to have generated a larger loss of social surplus. Thus, the lower the belief $\tilde{\varepsilon}_t$, the higher a and the lower L, the larger the penalty. The converse applies for the case $\tilde{\varepsilon}_t > \hat{\varepsilon}$.

The question the public are asking is: how good are the past actions of the decision maker? This is in contrast to a rational model, where the reputation would be a belief about the type of the decision maker, be it well informed or well intentioned.

I model the DM as a perfectly informed, fully rational agent, who knows the structure of the game, including the evaluation rule he faces. He maximises the expected discounted sum of reputation payoffs for the remainder of the game at each stage. I assume he knows the level of risk perfectly[10]. Let the discount rate of the DM be $\beta \in [0,1]$.

The game proceeds as follows. First, the DM chooses and implements the current period policy, y_t. If the activity is allowed, all observe the event z_t and all update their beliefs about risk. After this, the public evaluate the DM and he receives his reputation payoff. These stages are repeated in every subsequent period.

[10]It is a straightforward extension to allow him to have uncertainty over the level of risk. Specifically, as will be shown, the model will be linear so that he only cares about the expected value of ε at any time.

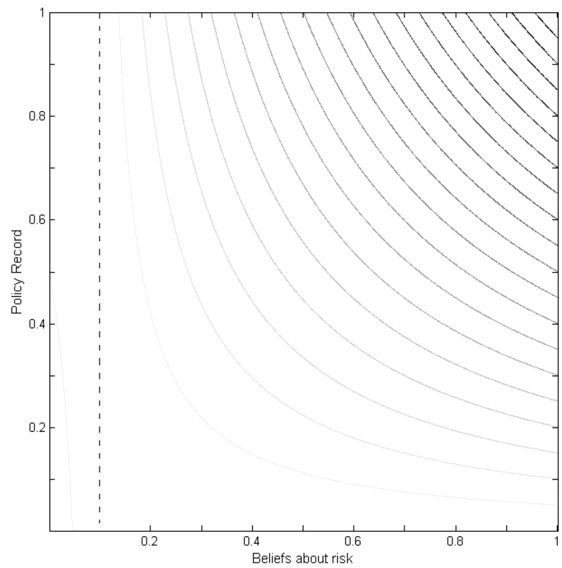

Figure 2.1: The reputation payoff function

2.4 Analysis of Baseline Setup

One can describe the problem in terms of a state space characterised by the belief of the public about risk and by the record of policy of the DM: $(\tilde{\varepsilon}, \tilde{y})$. The reputation payoff function can be mapped in this space. Figure 2.1 shows the contours of the reputation function when $\hat{\varepsilon} = 0.1$.

A dashed line separates two regions corresponding to the cases of low beliefs ($\tilde{\varepsilon} \leq \hat{\varepsilon}$), where there is public support for allowing, and high beliefs ($\tilde{\varepsilon} \geq \hat{\varepsilon}$), where the public want to ban the activity. Darker contour lines correspond to a lower reputation payoff.

Within each of these regions, the reputation payoff function varies monotonically. In the region of $\tilde{\varepsilon} \leq \hat{\varepsilon}$, the payoff increases as one moves northeast in the state space. In the region of $\tilde{\varepsilon} \geq \hat{\varepsilon}$, the payoff increases as one moves southwest in the state space.

When $\tilde{\varepsilon} \leq \hat{\varepsilon}$, the public support allowing the activity. The DM will be better off the higher his policy record is, and if his policy record isn't perfect ($\tilde{y} < 1$), then he will be penalised less when the public are close to indifferent than when they think the risk was extremely low and allowing had a higher expected payoff.

When $\tilde{\varepsilon} \geq \hat{\varepsilon}$, the public support a ban. The DM will have a perfect record if $\tilde{y} = 0$. If he has $\tilde{y} > 0$, then the higher the record (the more often/recently he has implemented the policy of allowing), the worse his reputation. Further, for a given record, the higher the perceived risk, the more he is penalised for having allowed.

On the dashed line itself, the payoff attains its highest possible value, zero, because the public are indifferent between the two actions so consider the DM has made no mistake, regardless of his past policy choices.

I make the following assumption for the analysis.

Assumption 1. (Small probabilities / volatile beliefs) For any initial belief, the observation of a single disaster event will cause the public to believe that banning is optimal: $1 - \delta \geq \hat{\varepsilon}$.

It will be useful to first outline some general properties of the model, for any number of periods. I refer to the expected discounted sum of reputation payoffs for some strategy as the value of that strategy. Since the reputation payoff has a maximum value of zero in each period, the maximum possible value of a strategy at any time is zero.

Lemma 5. *Whenever the public are indifferent between the two actions at the start of a period, $\tilde{\varepsilon}_{t-1} = \hat{\varepsilon}$, the DM can guarantee the highest possible value by banning for the rest of the game .*

Since the reputation payoff works as a penalty function and a public who are indifferent sees neither action as deserving of a penalty, banning will not be penalised under current beliefs about risk. Further, banning freezes beliefs and therefore ensures that initial beliefs are retained until the time of the evaluation. This logic holds no matter how many periods are left in the game.

Lemma 6. *If the public are in favour of banning, $\tilde{\varepsilon}_{t-1} > \hat{\varepsilon}$, and the DM has a record fully aligned with banning ($\tilde{y}_{t-1} = 0$), he can guarantee the highest value possible by banning for the rest of the game.*

Banning is able to achieve both goals of retaining current beliefs and perfectly aligning the record with current beliefs.

Lemma 7. *As long as there is risk, there is conservatism: that is, for any $\tilde{y}_{t-1} \in [0, 1]$, as long as $\varepsilon > 0$ and $\delta < 1$, there exist beliefs $\tilde{\varepsilon}_{t-1} < \hat{\varepsilon}$ such that it is optimal to ban.*

Proof. Lemma 5 says that starting from a state on the indifference line and banning forever, the DM can achieve the maximum possible value, zero. The payoff from banning increases continuously with respect to $\tilde{\varepsilon}_{t-1}$ in the region $\tilde{\varepsilon}_{t-1} < \hat{\varepsilon}$. Therefore, the value of the strategy "ban forever" must also fall continuously from zero as we move to initial states with $\tilde{\varepsilon}_{t-1} < \hat{\varepsilon}$. So one can make the value of such a strategy arbitrarily small by taking $\tilde{\varepsilon}_{t-1}$ close enough to $\hat{\varepsilon}$.

Instead, if the DM allows today, at least one of the outcomes will result in a reputation payoff that is strictly negative. Specifically, given Assumption 1, if there is a disaster, the payoff will certainly be negative at least for that period. Therefore, for any $(\tilde{\varepsilon}_{t-1}, \tilde{y}_{t-1})$, the value of a strategy that starts with $y_t = 1$ is at most $\varepsilon(\rho\tilde{y}_{t-1} + 1 - \rho)(a - (\delta\tilde{\varepsilon}_{t-1} + 1 - \delta)L) < 0$. This cannot be arbitrarily small.

The value of banning today is at least that from banning for all periods from today. Hence, there must be states close enough to indifference where the value of banning is higher than that from allowing initially. $\square$

Lemma 7 says that there will always exist states where the public would like for the activity to be allowed, but the DM will ban. I will refer to this phenomenon as conservatism - it is a case of the DM being *perceived* to be overly conservative; he may in fact be taking the socially optimal action.

In the remainder of this section, I start by considering the one-shot game. Since the beliefs are updated only after the action is chosen by the DM, even the one-period model already captures the essential trade-off for the DM. I then proceed to analyse a two-period model in which the DM can start with a clean slate and choose his initial record. This is followed by an analytical example of the infinite horizon game, where I restrict the DM's choice to a decision in the first period that will be locked in forever after. Finally, I use a numerical approach to confirm that results carry through to the general infinite horizon environment.

2.4.1 One-Period Model

Suppose the DM starts the game with some initial record $\tilde{y}_0$, aware that the public currently believe the risk of a disaster is $\tilde{\varepsilon}_0$, but that these will be updated after he sets policy and events are observed. What is the optimal choice of the DM for each initial state $(\tilde{\varepsilon}_0, \tilde{y}_0)$?

Using the feature of regional monotonicities, it is possible to show that the net incentive to ban - the difference between the expected payoff from banning and from allowing - is also monotonic with respect to the state variables, within each region. In addition, this incentive is continuous with respect to the components of the state. Therefore, the optimal strategy of the DM will be characterised by a threshold condition and these thresholds will vary continuously and monotonically as we move along one dimension of the state space within certain regions.

Proposition 3 describes the solution to the DM's one-period problem. It characterises the set of states where allowing and banning are optimal and describing the location of these sets and the boundaries separating them.

Proposition 3. *The optimal strategy of the one-period game is characterised by a partitioning of the state space into, at most, three sets. (i) There always exists a set of states where banning is optimal, which contains the states* $(\hat{\varepsilon}, \tilde{y}_0 \in [0,1])$ *and* $(\tilde{\varepsilon}_0 \geq \hat{\varepsilon}, 0)$. *(ii) There may or may not exist a (path-connected) subset of the region with* $\tilde{\varepsilon}_0 < \hat{\varepsilon}$ *where allowing is optimal and, when it exists, this region always contains the state* $(0,0)$. *(iii) There may or may not exist a (path-connected) subset of the region with* $\tilde{\varepsilon}_0 > \hat{\varepsilon}$ *where allowing is optimal and, when it exists, this region always contains a set of states with* $\tilde{y}_0 = 1$. *These subsets, call them* B, A_L *and* A_H, *can be described in terms of thresholds.*

- *For* $\tilde{\varepsilon}_0 \leq \hat{\varepsilon}$, *there exists a single threshold,* $\mu_L(\tilde{y}_0)$, *such that*

$$A_L = \{(\tilde{\varepsilon}_0, \tilde{y}_0) \,|\tilde{\varepsilon}_0 \in [0, \mu_L(\tilde{y}_0)), \tilde{y}_0 \in [0,1]\}$$

  if $\mu_L(\tilde{y}_0) \geq 0$ *for some* $\tilde{y}_0$, *and* $A_L = \emptyset$ *otherwise. There exists a level of true risk,* ξ, *such that if* $\varepsilon < \xi$, $\mu_L'(\tilde{y}_0) > 0$, *and if* $\varepsilon \geq \xi$, $\mu_L'(\tilde{y}_0) < 0$.

- *For* $\tilde{\varepsilon}_0 \geq \hat{\varepsilon}$, *two thresholds exist* $\mu_H(\tilde{y}_0)$ *and* $\mu_{HH}(\tilde{y}_0)$ *such that*

$$A_H = \{(\tilde{\varepsilon}_0, \tilde{y}_0) \,|\tilde{\varepsilon}_0 \in (\mu_H(\tilde{y}_0), \mu_{HH}(\tilde{y}_0)], \tilde{y}_0 \in [0,1]\}$$

  if $\mu_H(\tilde{y}_0) \leq 1$ *for some* $\tilde{y}_0$, *and* $A_H = \emptyset$ *otherwise. If* $\varepsilon \geq \hat{\varepsilon}$, *then* $\mu_{HH}(\tilde{y}_0) = 1$ *for all* $\tilde{y}_0$. *If* $\varepsilon < \hat{\varepsilon}$, *then there exist a subset of* $\tilde{y}_0$ *in* $\left[0, \frac{(1-\rho)\delta}{\rho(1-\delta)}\right]$ *for which* $\mu_{HH}(\tilde{y}_0) \in [\mu_H(\tilde{y}_0), 1]$. *For all parameter values,* $\mu_H'(\tilde{y}_0) < 0$ *and* $\mu_{HH}'(\tilde{y}_0) > 0$.

Proof in the Appendix.

I provide the graphed solutions for some example parametrisations in Figure 2.2. In all cases $a = 1$, $L = 10$ (implying a critical value $\hat{\varepsilon} = 0.1$), allowing the other parameters to vary: ρ, δ, and ε. Regions where the decision maker allows are shown in grey and those where he bans are shown in white.

Figures 2.2a and 2.2b show the solution when beliefs are more volatile than the policy record, with a lower actual risk on the left ($\varepsilon = 0.05 < \hat{\varepsilon}$) than the right ($\varepsilon = 0.2 > \hat{\varepsilon}$). Figures 2.2c and 2.2d are the solutions when the values of the parameters ρ and δ are exchanged and the policy record is more easily adjusted than beliefs.

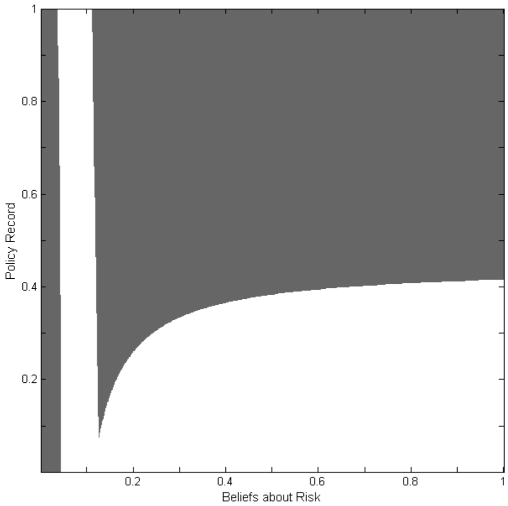

(a) $\rho = 0.9$, $\varepsilon = 0.05$, $\delta = 0.8$

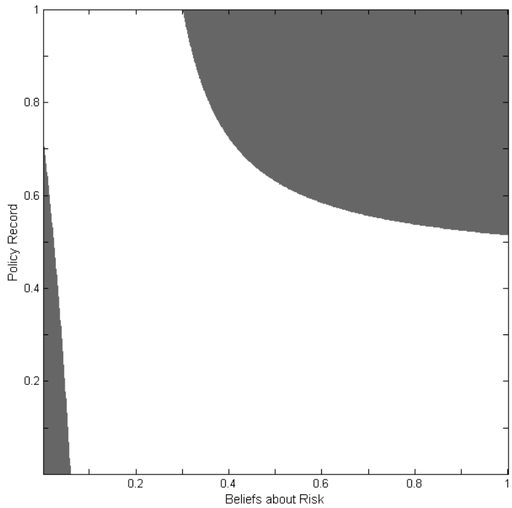

(b) $\rho = 0.9$, $\varepsilon = 0.2$, $\delta = 0.8$

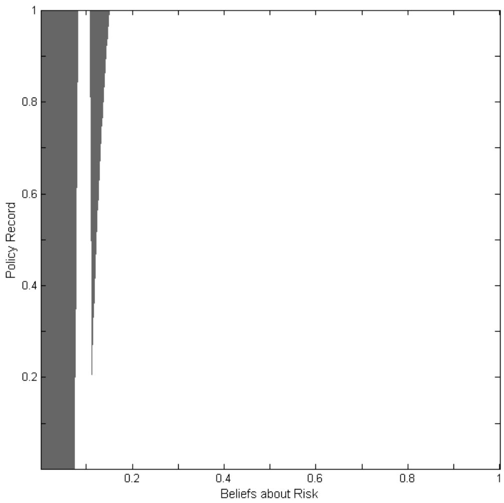

(c) $\rho = 0.8$, $\varepsilon = 0.05$, $\delta = 0.9$

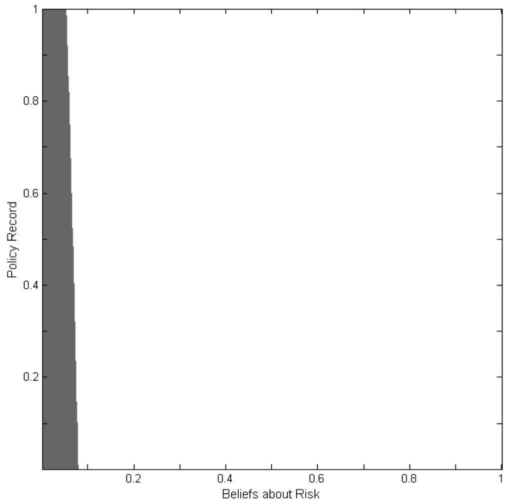

(d) $\rho = 0.8$, $\varepsilon = 0.2$, $\delta = 0.9$

Figure 2.2: Optimal decisions for a career concerned regulator

Dynamic incentives

The examples provided all confirm Lemma 7: for states with $\tilde{\varepsilon}_0 \lesssim \hat{\varepsilon}$, where the public are nearly indifferent between the two policies, banning will be preferred because the penalty for conservatism is low.

In the region of low perceived risk, $\tilde{\varepsilon}_0 < \hat{\varepsilon}$, the public currently believe that the DM should allow. Suppose the DM has a strong record in favour of allowing ($\tilde{y}_0 \approx 1$), his current situation is favourable. If he continues to allow, conditional on no disaster, he will look good tomorrow. But there is the risk that a disaster will happen and his high record will be heavily penalised. This is the case of hindsight bias: the public, who now think the risk is high, are sure they knew it all along, and so punish him for the act of allowing, even though at the time of allowing the public supported this action. If he bans, beliefs will be preserved but the record revised down, leading to some penalty.

It will be optimal to allow if the reputational risk is low (ε low, and δ high) and the reputational (sure) loss from eroding the policy record is high (ρ low). This explains why in the top left-hand corner of Figure 2.2a the DM allows but he doesn't in the same states of Figure 2.2b where risk is higher because of the higher ε. It also explains why there is more allowing in that corner in Figure 2.2c than in Figure 2.2a, because of the less volatile beliefs and more malleable policy record.

If his record of policy is strongly against allowing ($\tilde{y}_0 \approx 0$), banning the activity means his record will seem badly at odds with what he should be doing. By allowing, he better aligns his record with beliefs, which are likely to be preserved. But the DM also faces the risk that disaster will be observed, pushing beliefs up above the critical level. If his record doesn't adjust too much by the instance of allowing, then this swing in beliefs will make the still low record look not too bad. If the upwards revision of the record is large, the penalty when beliefs swing will be high.

Allowing will be optimal if the beliefs are likely to be preserved (ε small) and pandering is effective (ρ low), or if beliefs are likely to swing (ε high) and the record mostly preserved (ρ high). This explains why in the bottom left-hand corner, allowing happens for a larger set of states in Figure 2.2c compared to Figure 2.2a. In the former case pandering is more effective due to the lower ρ.

How the threshold varies with $\tilde{y}_0$ depends on the actual level of risk ε. If risk is actually extremely low, then the incentive to ban decreases with $\tilde{y}_0$ because at higher $\tilde{y}_0$ there are larger downward revisions to the record after banning, making conservatism less attractive ($\mu'_L(\tilde{y}_0) > 0$). If risk is not so low then conservatism is more attractive as $\tilde{y}_0$ increases because the penalty in the case of disaster would be particularly high ($\mu_L'(\tilde{y}_0) < 0$). This inversion of the slope of the boundary of region A_L can be seen comparing Figures 2.2c and 2.2d.

Next, let me consider the region where perceived risk is high, $\tilde{\varepsilon}_0 \geq \hat{\varepsilon}$. Here the public currently believe that the activity should be banned. Now banning has the double advantage of being the safe action and being able to align the record better with current beliefs. Therefore, at low $\tilde{y}_0$ it will be optimal to ban. Also, again close enough to indifference, it will be optimal to ban and avoid any swings in beliefs, since the penalty for $\tilde{y}_0 > 0$ is small.

When $\tilde{y}_0$ is high the DM faces a high penalty as things stand. If ρ is low and by banning he can align his record with beliefs quickly, he will do so (see Figures 2.2c and 2.2d, where he mostly bans when beliefs are high). But if this is not possible he can attempt to moderate beliefs. This requires allowing, so that the public will observe new events, most likely no disaster, and adjust down their belief about risk.

The details of the boundary of region A_H are subtle. First, when ε is small and $\delta < 1$, there exist some states close to, but not immediately adjacent to, the indifference line, where it is optimal to allow. Beliefs here can be reversed or at least brought down very close to indifference. It is worthwhile allowing because with high probability, allowing will look good under the new beliefs.

When the risk of disaster is small but not too small, the incentive to try to moderate beliefs is strongest when beliefs are most extreme because the marginal damage of disaster is smaller: beliefs cannot get much worse. Allowing will happen in the top right-hand corner of the state space. This is what we see in Figure 2.2b. Instead when the risk of disaster is very small the dominant trade-off is between the no disaster scenario, where beliefs are moderated but the record pushed up, and banning, where beliefs remain extreme but the record is better aligned. For low initial record, $\tilde{y}_0$, banning is quite attractive because the penalty won't be too large. While allowing can be attractive when one can moderate beliefs close to indifference, this becomes too costly with more extreme beliefs. This explains the case in which the boundary between the banning and allowing regions slopes up in some region and we get the two thresholds. We observe this in Figures 2.2a and 2.2c.

Generally, both actions may align the record better with beliefs, depending on whether beliefs are above or below the critical level. But the two actions have very different effects on learning. Banning has the advantage of being a "safe action" because there is no uncertainty over future payoffs, no large swings in beliefs are possible. In contrast, allowing reveals new information which may shift beliefs to be more or less favourable for the DM.

Comparative Statics

The discussion of the dynamic incentives already suggests how the sets B, A_L and A_H will vary with the key parameters of the problem, δ and ρ, since these determine the elasticity of beliefs about risk and about the record of policy.

Proposition 4. *The thresholds have the following comparative statics:*

- $\frac{\partial \mu_L}{\partial \delta} > 0$ *and, for every* $\tilde{y}_0$*, there exists a value of actual risk,* ξ*, such that if* $\varepsilon < \xi$ *then* $\frac{\partial \mu_L}{\partial \rho} < 0$ *and if* $\varepsilon \geq \xi$ *then* $\frac{\partial \mu_L}{\partial \rho} \geq 0$*, where* ξ *is increasing with* $\tilde{y}_0$ *and* $\xi > 1$ *at* $\tilde{y}_0 = 1$*;*
- $\frac{\partial \mu_H}{\partial \delta} < 0$ *when* $\mu_H < \frac{\hat{\varepsilon}}{\delta}$*,* $\frac{\partial \mu_H}{\partial \delta} > 0$ *when* $\mu_H > \frac{\hat{\varepsilon}}{\delta}$ *and* $\frac{\partial \mu_H}{\partial \rho} < 0$*;*
- $\frac{\partial \mu_{HH}}{\partial \delta} < 0$ *and* $\frac{\partial \mu_{HH}}{\partial \rho} > 0$*.*

Proof in the Appendix.

For μ_L, higher δ (more stable beliefs) reduces the risk of allowing because the beliefs after a disaster won't be too extreme, therefore there will be less conservatism. The effect of ρ depends on actual risk. Consider the case of $\tilde{y}_0 = 0$; ρ determines how much the record will adjust if the DM allows. If risk is low enough, then allowing is less attractive the higher ρ because the DM cannot pander to current beliefs effectively. Instead if risk is relatively high, allowing is more attractive if the record doesn't adjust too much, because when disaster occurs, a low record will look better. Now consider $\tilde{y}_0 = 1$, in this case ρ affects how much the record will adjust when the DM bans. Conservatism is more costly if the record will be eroded away a lot by one instance of banning, so the threshold will decrease with ρ regardless of the riskiness of allowing (i.e. regardless of ε).

What happens in the region with high beliefs, $\tilde{\varepsilon}_0 > \hat{\varepsilon}$? Close to the indifference line, when beliefs are more volatile (lower δ), after the public observe no disaster the belief is moderated to a lower level, beyond the indifference level. But the reputation payoff will be highest close to indifference, therefore there will be a stronger incentive to allow and alter beliefs when these can be brought as close as possible to indifference, i.e. when δ is higher. Instead, for initial beliefs higher than $\frac{\hat{\varepsilon}}{\delta}$, it is more tempting to allow and try to moderate beliefs, the more these can be lowered after observing no disaster. Hence higher δ will lead to more banning here. High ρ will increase the incentive to allow, and therefore expand the A_H area, because banning now is ineffective at moderating the policy record and the DM is more willing to take a risk to bring down perceived risk.

Where things go wrong and where things appear to go wrong

It is clear that there is plenty of scope for inefficient behaviour. I need to consider two things: when is the DM taking an action that is inefficient according to the true risk, and when is he doing so from the point of view of the public? This would correspond to an objective and subjective definition of misbehaviour.

Remark. If the risk is low ($\varepsilon \to 0$), there will always exist states where the DM is objectively inefficient. If the risk is high enough ($\varepsilon \to 1$), it is generally the case that the DM is efficient in every state.

The strong bias towards the safe action, as described by Lemmas 5 and 6, implies that banning must take place in some non-empty set of states. This is true even when the risk goes all the way to zero[11]. In contrast, high enough risk can dissuade a DM from allowing at any state.

Remark. The optimal action of the DM may contradict both the objectively correct policy and the perceived correct policy, and do so simultaneously, for either correct policy.

Perhaps the most surprising case is that in which the DM chooses to allow the activity even though the public are strongly opposed to this, and he knows risk to be above the critical level. It happens as a result of the inability of the DM to clear his policy record. If he is heavily invested in his position of supporting deregulation, and a single observation of no disaster can effectively moderate beliefs, he may have more to gain from allowing than banning, even though it is socially undesirable.

2.4.2 Two-Period Model

I will make the following assumption in the two period model: in the first period $\rho = 0$, so that the first period choice fully determines the record of policy for that period. At the start of period 2 the policy record will be extreme, $\tilde{y}_1 \in \{0, 1\}$.

Proposition 5. *In the two period game, the optimal action by the DM in each period $t \in \{1, 2\}$ is characterised by a set of thresholds. In period 1, there exists a $\mu_{L,1} < \hat{\varepsilon}$ such that it is optimal to allow if and only if $\tilde{\varepsilon}_0 < \mu_{L,1}$. In period 2, if $\tilde{y}_1 = 0$, there exists a $\mu_{L,2}(0) < \hat{\varepsilon}$ such that it is optimal to allow if and only if $\tilde{\varepsilon}_0 < \mu_{L,2}(0)$; if $\tilde{y}_1 = 1$, then for the case that $\rho \geq \delta$ there exist two threholds, $\mu_{L,2}(1) < \hat{\varepsilon}$ and $\mu_{H,2}(1) > \hat{\varepsilon}$, such that it is optimal to allow if and only if $\tilde{\varepsilon}_0 < \mu_{L,2}(1)$ or*

[11] With the exception of the special parameter case, $\delta \leq \hat{\varepsilon}$ and $\rho = 0$, where the DM can moderate even the most extreme beliefs to below the critical value and simultaneously perfectly align his record with allowing.

$\tilde{\varepsilon}_0 > \mu_{H,2}(1)$*, and for the case that* $\rho < \delta$*, there exists an additional threshold* $\mu_{HH,2}(1) > \mu_{H,2}(1)$ *such that it is optimal to allow if and only if* $\tilde{\varepsilon}_0 < \mu_{L,2}(1)$ *or* $\tilde{\varepsilon}_0 \in (\mu_{H,2}(1), \mu_{HH,2}(1))$.

Proof in the Appendix.

The qualitative solution is the same as in the one period model. The additional insight is that, if allowed to choose one's initial record, then the DM will ban in any state where perceived risk is above the critical level. The evaluation procedure provides no incentive for a DM to educate the public if they are mistaken in thinking the activity too risky: the public will not be suspicious of a DM that acts according to their prior on risk, and allowing always entails a risk to the DM.

Setting $\rho = 0$ in the first period means we can focus on the top and bottom edges of the state space for period 2, i.e. $\tilde{y}_1 = 0$ and $\tilde{y}_1 = 1$. Supposing the DM originally banned, the picture looks similar, banning is the optimal action except possibly for a range of beliefs close to $\tilde{\varepsilon}_1 = 0$. Allowing later on in the game may be less attractive because the DM has no opportunity to fix things if disaster occurs, however, it also reduces the opportunities for a disaster to be observed.

If the DM allowed in the first period, there may be two regions of beliefs where he finds it optimal to persevere and continue to allow in the second period. This happens again at the lowest $\tilde{\varepsilon}_1$, where the public would like him to allow and he can pander to their beliefs, and possibly at a set of high $\tilde{\varepsilon}_1$, where he is facing a high penalty and may prefer to address it by moderating beliefs rather than adjusting the policy record.

2.4.3 Irreversible Policies

Extending the framework to infinite periods is not trivial despite the recursive structure of the problem. To make the problem analytically tractable I focus on the case of irreversible policies. This is also an interesting application because many regulatory decisions are to some degree costly to reverse. This also allows me to check whether the lack of penalties for changing policy may have played an important part in shaping the results presented so far.

Even with the restriction of irreversible policies, solving the problem analytically is complicated by the kink in the reputation function because one needs to know what side of $\hat{\varepsilon}$ the belief is on, for every history of shocks $\{z_1, z_2, \ldots\}$. I therefore analyse the problem for the extreme case of $\delta = 0$ where beliefs are maximally volatile. For this case, it is possible to write down the value of each action at each initial state.

I show that the decisions will depend on current perceptions of risk and on the DM's initial stance on regulation in exactly the same way as the previous finite time models with reversible policies.

Proposition 6. *In the case of irreversible policies, and with* $\delta = 0$*, the optimal strategy is characterised by a partitioning of the state space into, at most, three sets. (i) There always exists a set of states where banning is optimal, which contains the states* $(\hat{\varepsilon}, \tilde{y}_0 \in [0,1])$ *and* $(\tilde{\varepsilon}_0 \geq \hat{\varepsilon}, 0)$*. (ii) There may or may not exist a (path-connected) subset of the region with* $\tilde{\varepsilon}_0 \leq \hat{\varepsilon}$ *where allowing is optimal; this region, when it exists, always contains the state* $(0,0)$*. (iii) There may or may not exist a (path-connected) subset of the region with* $\tilde{\varepsilon}_0 \geq \hat{\varepsilon}$ *where allowing is optimal; this region, when it exists, always contains the state* $(1,1)$*. These subsets, call them* B*,* A_L *and* A_H*, can be described in terms of thresholds:*

- *A single threshold,* $\mu_L(\tilde{y}_0)$*, for* $\tilde{\varepsilon}_0 \leq \hat{\varepsilon}$ *such that*

$$A_L = \{(\tilde{\varepsilon}_0, \tilde{y}_0) \,|\tilde{\varepsilon}_0 \in [0, \mu_L(\tilde{y}_0)), \tilde{y}_0 \in [0,1]\}$$

  if $\mu_L(\tilde{y}_0) \geq 0$ *for some* $\tilde{y}_0$*, and* $A_L = \emptyset$ *otherwise. There exists a level of true risk,* ξ*, such that if* $\varepsilon < \xi$*,* $\mu_L'(\tilde{y}_0) > 0$*, and if* $\varepsilon \geq \xi$*,* $\mu_L'(\tilde{y}_0) < 0$*.*

- *A single threshold,* $\mu_H(\tilde{y}_0)$*, for* $\tilde{\varepsilon}_0 \leq \hat{\varepsilon}$ *such that*

$$A_H = \{(\tilde{\varepsilon}_0, \tilde{y}_0) \,|\tilde{\varepsilon}_0 \in (\mu_H(\tilde{y}_0), 1], \tilde{y}_0 \in [0,1]\}$$

  if $\mu_H(\tilde{y}_0) \leq 1$ *for some* $\tilde{y}_0$*, and* $A_H = \emptyset$ *otherwise. Whenever region* A_H *exists, we have* $\mu_H'(\tilde{y}_0) < 0$*.*

Proof in the Appendix.

The reason why there is no longer the case where the boundary of A_H has an upward sloping section is because of the assumption that beliefs are extreme ($\delta = 0$). Earlier we described how, if the beliefs were very extreme, and the DM was not confident that observations of no disaster would succeed in moderating them, he might give up and ban in order to align his policy record with the pessimistic beliefs. Now this is not a concern, and at higher beliefs the value of allowing is the same, while the value of banning decreases for any $\tilde{y}_0 > 0$; so the DM will ban for all beliefs above the threshold $\mu_H(\tilde{y}_0)$.

2.4.4 Infinite horizon, numerical solution

Solving the problem analytically for an infinite horizon problem is not possible because of the discrete jumps in the state variables - one cannot write down the expected discounted flow of payoffs along any given path. However, the problem is stationary and I can solve it numerically.

The Bellman equation for the problem is

$$V\left(\tilde{\varepsilon}_t,\tilde{y}_t\right)=max\left\{\begin{array}{c}\Delta\left(\tilde{\varepsilon}_t,\rho\tilde{y}_t\right)+\beta V\left(\tilde{\varepsilon}_t,\rho\tilde{y}_t\right),\\ \left(1-\varepsilon\right)\Delta\left(\delta\tilde{\varepsilon}_t,\rho\tilde{y}_t+1-\rho\right)+\varepsilon\Delta\left(\delta\tilde{\varepsilon}_t+1-\delta,\rho\tilde{y}_t+1-\rho\right)\\ +\beta\left[\left(1-\varepsilon\right)V\left(\delta\tilde{\varepsilon}_t,\rho\tilde{y}_t+1-\rho\right)+\varepsilon V\left(\delta\tilde{\varepsilon}_t+1-\delta,\rho\tilde{y}_t+1-\rho\right)\right]\end{array}\right\}$$

The operator on the RHS of the equation defines a mapping, call it T. T satisfies Blackwell's sufficient conditions for a contraction mapping (monotonicity and discounting). Therefore I can use the contraction mapping theorem: the solution to my problem is the fixed point of a contraction mapping and, further, I can iterate on the Bellman equation from any initial (continuous and bounded) guess for the value function and can be certain to converge to the solution.

This is precisely the algorithm that I use in the numerical approach, starting with the simplest candidate value function $V_0\left(\tilde{y}_0,\tilde{\varepsilon}_0\right)=0$ for all $\tilde{y}_0,\tilde{\varepsilon}_0\in[0,1]$. The only simplification required now is a discretisation of the state space, but since my action set is binary I am able to choose quite a fine grid: 1000×1000. Nevertheless, the discretisation has caused the boundaries of the region not to be perfectly smooth in some cases; the jaggedness of the boundaries is an artifact of the discretisation, not a new feature of the solution in the infinite period model.

The results are qualitatively equivalent to the one period case. Figures 2.3a-2.3d show the solutions to the problem using all the same parameters as in the examples we studied in the one period model (with one additional parameter, $\beta=0.7$).

The main conclusion of all of these simulations, and of all the earlier results, is that whether the regulator will ban or allow depends not only on the fundamentals of the process being regulated (the actual risk ε, and the gain and loss parameters). Instead the decisions will be strongly influenced by the psychological parameters, ρ and δ. The process by which perceptions of risk update and by which regulators are able to rebrand themselves as pro or anti regulation, will determine in large part how regulators behave. This decoupling of decisions from the fundamentals of the process will of course generate suboptimal policies.

The properties from the one period case carry through. First, it is clear that the benefits of the safe action for the decision maker still play an important role; there

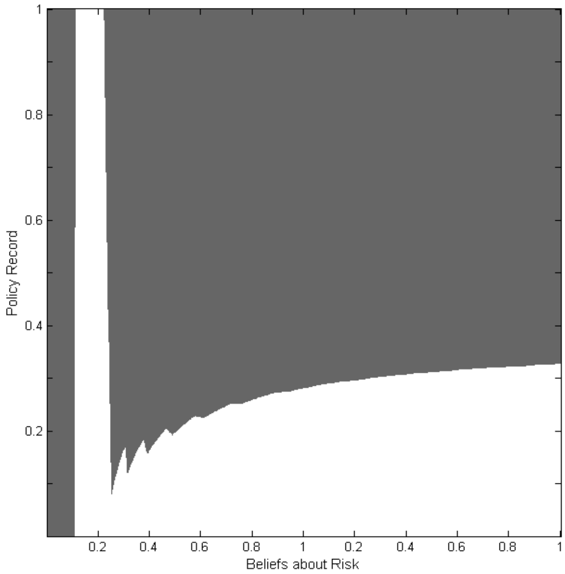

(a) $\rho = 0.9$, $\varepsilon = 0.05$, $\delta = 0.8$

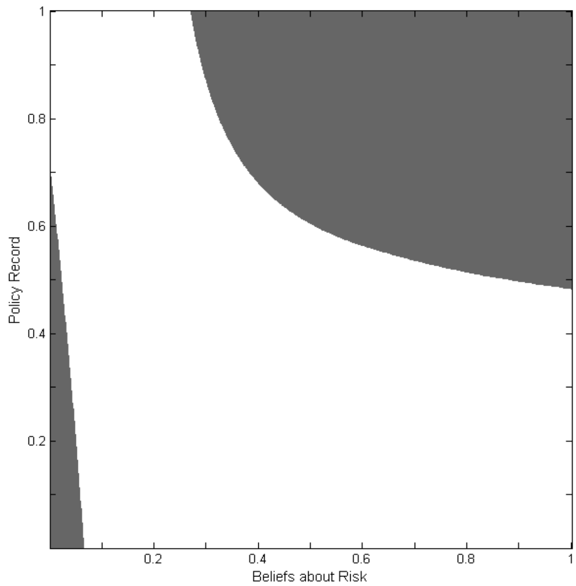

(b) $\rho = 0.9$, $\varepsilon = 0.2$, $\delta = 0.8$

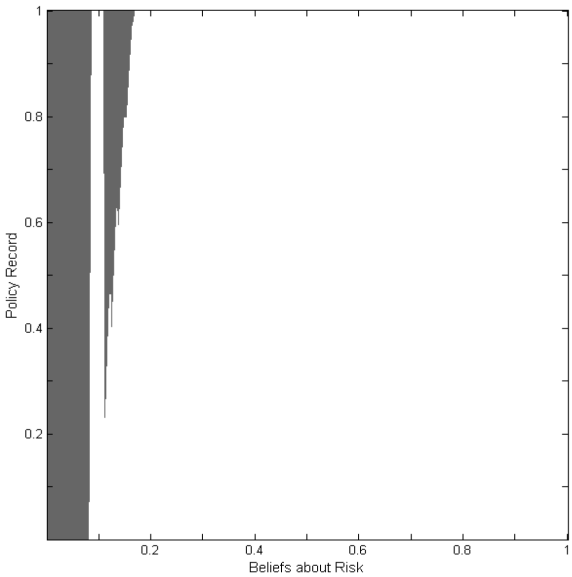

(c) $\rho = 0.8$, $\varepsilon = 0.05$, $\delta = 0.9$

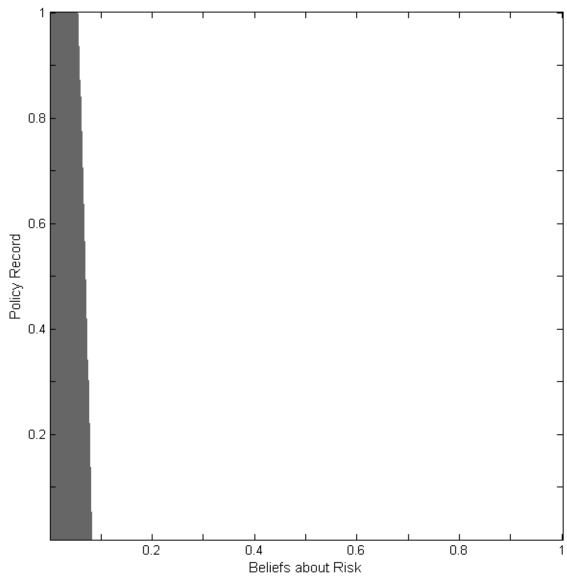

(d) $\rho = 0.8$, $\varepsilon = 0.2$, $\delta = 0.9$

Figure 2.3: Optimal decisions for a career concerned regulator - infinite horizon

is always banning close to indifference and when beliefs are above the critical level and the record is low. Allowing is optimal in regions which are essentially described by Proposition 3: the state $(0,0)$ is always included in A_L, A_H always contains high $\tilde{y}_0$ states. A higher actual level of risk always shrinks the set of allowing states.

Paths, events and absorbing states

There always exists a set of states where the DM bans forever once it reaches one of these states - absorbing states. There are four cases for the paths in which the decision maker allows at some point, which depend on the parameters of the problem and the initial beliefs.

1. It may be that if disaster occurs early on, the DM falls into a banning absorbing state, but that once he has allowed for a period of time, he will allow no matter what in the future. If actual risk is low enough then he will find it optimal to maintain a strong record of allowing, and when disasters occur he may be confident in his ability to moderate beliefs if he perseveres in allowing. This happens when ε is small, and $\rho > \delta$. These are the dynamics of the case in Figure 2.3a.

2. He may allow if the public continue to support the activity but switch to banning if disaster induces fear. When the beliefs are not too volatile and risk not too high he will find it optimal to continue to allow as long as the public are strongly in favour of it. After a disaster however, he may find it easier to change his record of policy by starting to ban. In this case then, for an infinite horizon model, allowing is partly steady, or conditionally steady (conditional on no disaster) and banning is a steady state. This happens when ε is not too high and $\rho < \delta$. Such dynamics arise in the cases of Figures 2.3c and 2.3d

3. He may become cautious while beliefs are in his favour, but be stubborn in allowing if beliefs turn against him. If the record of policy is quite difficult to adjust after the first action, then the DM prefers to ban if beliefs currently favour him as he can avoid swings in beliefs while retaining a reputation for being (mostly) in favour of allowing. If beliefs are against him, the best he can do is allow and moderate beliefs since it will be difficult to persuade the public he now supports a ban on the activity. What we observe then is that as long as beliefs are in favour of allowing, the DM will tend towards a moderate policy record; if beliefs turn against him he will start allowing. If he can persuade the public that the activity is quite harmless, then he will reenter the cycle that yielded a moderate record. If he cannot moderate beliefs enough then he

will allow for a few periods, until beliefs are less extreme and eventually enter the banning absorbing state. This will happen at intermediate values of ε and when $\rho > \delta$. This is what happend in Figure 2.3b

4. Finally he may become cautious in the future no matter what. If the risk is relatively high (from his point of view) he may always want to maintain a moderate record, and may not feel it worth the risk to continue to allow once a disaster has occurred. In that case eventually all paths end up in the banning absorbing state. This will happen at intermediate values of ε and when $\rho < \delta$.

When ε is high there will be no allowing.

2.5 Institutional variations

I consider how changes to the institutional constraints faced by the DM affect outcomes and, where possible, I comment on whether certain measures are welfare-enhancing.

Outside signals

Suppose now that the activity is being pursued in a neighbouring locality, and the public get to observe the outcomes there. Now the decision maker cannot completely control the beliefs of the public, in particular, he no longer has a safe action. I assume that if the public observe only the outside signal, s, which is identically and independently distributed as the local shock z, they adjust beliefs according to the rule

$$\tilde{\varepsilon}_t = \delta^{neigh}\tilde{\varepsilon}_{t-1} + \left(1 - \delta^{neigh}\right) s_t$$

and if both events, local and from outside, are observed then the adjustment will be

$$\tilde{\varepsilon}_t = \delta\left(\delta^{neigh}\tilde{\varepsilon}_{t-1} + \left(1 - \delta^{neigh}\right) s_t\right) + (1 - \delta)\, z_t$$

The solution is still of the same form as the earlier problem[12]. I focus on some examples to illustrate the basic point that increasing information (as long as the public factor it into their learning to some extent) generally improves policy outcomes, but not for every state and parameter configuration.

[12]The full solution to the problem will be a complicated function of the parameter δ^{neigh} and the initial state, because now whether in the next period there is a jump between regions (above and below $\hat{\varepsilon}$) depends on the initial state and the two events z_t and s_t so there are many different cases to consider.

First consider the region of low perceived risk, $\tilde{\varepsilon}_0 \leq \hat{\varepsilon}$. When actual risk is very low, the presence of outside signals expands the allowing set; when the risk is very high it shrinks the allowing set. For intermediate values of ε, the threshold with outside signals is higher at low $\tilde{y}_0$ and lower at high $\tilde{y}_0$.

The intuition is as follows. When the risk is extremely low, the outside signal will almost certainly result in $s_1 = 0$, making the new beliefs more extreme, $\tilde{\varepsilon}_1 < \tilde{\varepsilon}_0$. This makes it more costly to ban, because a lower record will be assessed against the lower $\tilde{\varepsilon}_1$, and therefore the incentive to allow increases. When the risk is high, the probability that at least one of the two processes will generate a disaster becomes substantial, and even that of two disasters is non-negligible; therefore there is a stronger incentive to ban in the presence of more signals. Finally, if the actual risk is small but not so small, the probability of two disasters is not really significant, but the probability of at least one disaster is. At low $\tilde{y}_0$ this increases the incentive to ban, because the DM hopes that an $s_1 = 1$ event will persuade the public that risk is high while he maintains a low record, yielding a very high payoff if disaster does occur. For high $\tilde{y}_0$ instead, there will be less conservatism. The s_1 signal serves as a form of insurance against extreme beliefs; the probability of observing one disaster is higher but it is counteracted by observing none for the other process and this reduces reputational risk overall.

Notably now, when risk is extremely low, there may be no conservatism at all at high $\tilde{y}_0$ so that the DM will allow even when the public are indifferent. This is the result of his inability to freeze beliefs. Since beliefs will be lowered by the s_1 signal, if ρ is relatively low, the DM will not want to erode his record when he anticipates beliefs will change in favour of allowing.

In the region of high perceived risk, $\tilde{\varepsilon}_0 \geq \hat{\varepsilon}$, let's consider a couple of particular cases. For example, first suppose that ρ is high ($\rho > \delta$) so that in the model without outside signals we know there is a strong incentive to allow in order to moderate beliefs. Further, suppose that the boundary between the banning and allowing regions lies far enough east that even if $z_1 = s_1 = 0$ was observed, beliefs would remain above the critical value (they can be moderated, but banning will still be favoured). This is the case for intermediate ε (e.g. case in Figure 2.3b). In this case, the presence of outside signals increases the incentive to ban and the set A_H shrinks. The reason is that the DM can now rely on the outside signal moderating beliefs, while he moderates the record. This can prevent situations of inefficient stubbornness, where the DM sees no option but to allow as a way to moderate public beliefs.

Finally consider the case of low actual risk (the case of Figures 2.3c and 2.3a) and δ is relatively high. In the model without outside signals it was optimal to ban close to indifference but also at more extreme beliefs. The reason why the DM was banning for $\tilde{\varepsilon}_t \geq \mu_{HH}(\tilde{y}_t)$, was that beliefs couldn't be sufficiently moderated; if the public are too pessimistic, the DM gives up and bans. The presence of outside signals will make allowing more attractive, because he can be quite certain that the public will observe two instances of no disaster and the beliefs can be brought down further. Again, more information can reduce inefficiency; in this case, excessive regulation.

Summary of effects of outside signals. Allowing for additional information in the form of external signals generated by an identical process can be helpful, especially for preventing excessive regulation when the risk is really very low. The additional signals avoid the problem of policy getting trapped in the (no longer) safe action. The public cannot be prevented from acquiring information and the DM cannot exploit a state of perpetual ignorance. Further, multiple signals reduce how volatile beliefs are and this reduces the reputational risk of allowing.

Finally, if the public's rule for learning did not have the feature of forgetfulness then allowing for additional signals would be strictly welfare improving in the long run because whether or not the DM himself provided information by allowing himself, the public would converge to the truth using the s signals. Eventually, the DM would find it harder to manipulate beliefs and would be limited to aligning his record with the true best policy.

Longer office terms

Next, I consider the effect of lengthening career terms. I do this by looking at the comparative static on β in the model I analysed in section 2.4.3 with irreversible policies. One can interpret β as the probability of staying in office for one period longer. Under this interpretation an increase in β represents longer average time in office. The effect of β on decisions is driven by the long run payoffs that are reached only if the DM stays in office long enough.

For the threshold in the region $\tilde{\varepsilon}_0 \geq \hat{\varepsilon}$ we have

$$sign\left(\frac{\partial \mu_H}{\partial \beta}\right) = sign\left(\frac{\varepsilon\,(1-\rho)}{\rho \tilde{y}_0\,(1-\beta)}\right)$$

The threshold in this region is thus increasing in the expected time in office. In other words, if the public are currently against allowing the activity, then lengthening career terms will make it more likely that the DM bans.

The long run value of banning is very high when perceived risk is high, call it

$V_0^\infty = 0$. Instead the long run value of allowing is $V_1^\infty = \varepsilon(a - L)$, which is strictly negative for any $\varepsilon > 0$. If the DM is more confident he will stay in office for long, then banning becomes more attractive because of the higher probability of reaching the value V_0^∞. Hence the threshold moves out, and the banning region B expands.

This feature carries over to the infinite horizon models without the additional assumption about irreversibility of policies. It is driven by the ability to generate perfect payoffs in the future (with certainty) by banning, which is not possible along paths where the DM allows.

Since this property is true for all values of ε, the tendency towards tighter regulation as terms get longer may be welfare-decreasing. For values of risk below the critical level we would like for the DM to take more risk and educate the public, but he will be less willing to do so when he has time to realign his policy record with public beliefs while facing no reputational risk.

For the threshold in the region of low perceived risk, $\tilde{\varepsilon}_0 \leq \hat{\varepsilon}$,

$$sign\left(\frac{\partial \mu_L}{\partial \beta}\right) = sign\left((1-\varepsilon)\,\hat{\varepsilon} - \varepsilon\,(1 - \rho\tilde{y}_0)\right)$$

Now the long run payoff from banning is not zero, it is $V_0^\infty = -\left(a - \tilde{\varepsilon}_0 L\right) < 0$. The long run payoff from allowing is the same as before.

For $\varepsilon \to 0$, the threshold is increasing with the average time in office because V_1^∞ is close to zero. So in this case, conservatism will decrease. For large enough ε, V_1^∞ will be very low and the threshold will decrease with β. Therefore, in the limits $\varepsilon \to 0$ and $\varepsilon \to 1$, longer office terms generate better policies: low actual risk leads to more allowing and high risk to more conservatism.

The effect of β also depends on how heavily invested the DM is in his regulatory position. At low $\tilde{y}_0$ longer terms generally lead to better policy choice. The condition at $\tilde{y}_0 = 0$ becomes

$$sign\left(\frac{\partial \mu_L}{\partial \beta}\right) = sign\left((1-\varepsilon)\,\hat{\varepsilon} - \varepsilon\right)$$

For small ε probabilities, this is not too different from $sign\left(\hat{\varepsilon} - \varepsilon\right)$, so that the effect of higher β on the threshold is welfare increasing.

However, it may well be the case that for $\varepsilon > \hat{\varepsilon}$ we get more pandering in the region of low beliefs if the probability of survival is higher, especially at high $\tilde{y}_0$. This will hinge on the parameter ρ: when ρ is low, then the record quickly adjusts down after banning and the costs of choosing the banning path (given that $\tilde{\varepsilon}_0 < \hat{\varepsilon}$) are experienced soon, so that β doesn't much affect V_0. When ρ is high instead, these costs arrive only after several periods so higher β yields a lower V_0 and a stronger incentive to pander. This means that the threshold μ_L may be increasing in β at values of ε well above the critical level. In those instances, longer terms in office make things worse.

Summary of the effects of longer office terms. Increasing the length of terms in office will exacerbate the bias toward tighter regulation when public perceptions of risk are high, often to the detriment of public welfare. In the region of low perceived risk, extending terms is a good measure in the limits of low and high risk and when dealing with near zero probabilities, but it may exacerbate the problem of inefficient pandering for intermediate probabilities.

Policy flexibility

Finally, I evaluate the effect of making it costlier for the DM to change policy. Consider, as an example, the dynamic paths in Figure 2.2d. For any path where the decision maker would not change course, introducing a cost to switches will make no difference. So crucially, in the absorbing states where the DM bans forever, adding a cost to switching will not affect outcomes. In what follows, I consider the cases where the policy maker will want to change course with some probability.

Consider region A_H where perceived risk is high, but the DM would like to allow. He does so because his policy record is high and this is hard to reverse. However, if the perceived risk is extremely high, he is not hopeful that he will be able to convince the public that allowing is optimal. Instead, he will allow until beliefs are moderated somewhat, and then he will start banning. On the boundary of the region, he is indifferent between giving up now and banning, or attempting to moderate beliefs. When we make switching costly, the value of banning forever doesn't change, but the value of allowing does because it will then be costly to switch to banning later. Therefore, banning in these states must be more likely ex-ante with costly switching. In the states further from the boundary, it will still be worthwhile for the DM to allow, and ex-post banning will be reduced.

In the region of low perceived risk, consider the states with low $\tilde{y}_0$ in the A_L region. Here the DM is allowing, but since his record is low, if a disaster does take place, he would very much like to switch to banning at this point. Again, removing the option to switch (or making it costly) will reduce the long-run value of allowing and lead to more banning ex-ante (and less banning ex-post). This is the same pattern as for the A_H region. In the region of higher $\tilde{y}_0$ we sometimes find that the DM tries to maintain a moderate policy record; in the case of binary actions, this involves cycling between allowing and banning (conditional on no disaster). But the intuition described above carries through: allowing still loses value when the DM will not be able to switch to banning after disaster, and therefore the DM will be more likely to ban ex-ante.

Summary of the effects of making policy switches costly. Making it costly (or impossible) for the DM to switch affects decisions through two channels. Ex-post the DM will be forced to persevere with policies that he may wish to abandon, and this can reduce inefficient over-regulation, especially after disasters are observed. But, anticipating this, the DM will be less willing to allow ex-ante compared to the case of flexible policies.

2.6 Comparison of Full Rationality and Bounded Rationality Predictions

There is quite a lot of freedom in how one could specify a fully rational model for this environment. Crucially we need to specify more about the incompleteness of information: we need a notion of quality of decision makers, a set of possible states of the world (risk levels), and some beliefs over the distribution of qualities and risk levels at the start of time. The most natural case is for policy makers to differ in their information about the risk: suppose good types know the level of risk perfectly and bad types have a noisy signal.

In a one-period version of such a model, one result is that a decision maker that chooses to allow and is successful (i.e. a disaster is not observed) always has the highest reputation regardless of the priors. Suppose that the bad type gets a low risk signal with the same probability as a good type, but it is more weakly correlated with the true state. If both types were to follow their signals, then the bad type should more often pick to allow when the risk was in fact high, so that observing no disaster should cause the posterior to be higher than the prior, and observing a disaster should cause a lower posterior. In fact, there is no pooling equilibrium (where both follow their signals) because the bad type would do worse from allowing, even if he observed a low-risk signal. The equilibrium will be semi-separating: the bad type will ban after a high-risk signal, but after a good signal he will also ban with some probability because banning has a lower risk. This again pushes up the posterior after allowing and no disaster. Taking the risk and succeeding always receives the highest reward.

If the probability of disaster is very low, then Bayesian updating will result in a no-disaster event making very little difference to beliefs. That is, the reputation after allowing and no disaster is very close to that after banning; voters are generally ambivalent about policy makers (the exception being when disaster happens; this does significant damage to their reputation). Another way of saying this is that, in equilibrium, the public reward dissenting decision makers (who go against the prior) as much as those who conform to prior public beliefs (conditional on no disaster).

In the rational model the most informed types always follow their signal. The arbitrage of the two actions will make the less informed types indifferent, and the good types (with their superior information) strictly prefer the right action. Therefore distortions are limited, the bad types will ignore their signal with some probability by playing the action that is less often induced by their information. These features are true even if good types are not perfectly informed and for any prior in the one period game.

Next consider a repeated game, for simplicity, suppose there are two periods. The state of the world is realised at the start of the game and stays constant in all future periods: risk is the same over time. Assume that a policy maker obtains all his information at the start of time; later on in the game, he only observes the outcomes of allowing, the same as the public. Suppose that the good type is perfectly informed, or has sufficiently accurate signals that observing disaster does not lead him to reverse his belief that the true state of risk is low after a low risk signal.

In equilibrium, the good types should follow their signal at all times, but then the bad types must also choose constant policies, otherwise they would reveal themselves. So an equilibrium exists in which out of equilibrium beliefs are that any decision maker deviating in their policy is a bad type and no one deviates. Such a model cannot explain paths in which regulators change policies over time, and in response to events, something that is observed often in the real world.

Are there equilibria where deviation is possible? If the good type has a perfect signal, this cannot be the case. Since the public know that good types follow their signal in the first period, they will never reward deviation. What if the good type has an imperfect signal? Then as long as his signal is accurate enough that a single disaster event yields a posterior that still favours allowing, the same is true. The good type should continue to allow; and therefore the bad type will mimic him and persevere in his policy choices too.

The driver of distortions in this model is that the bad type is not able to correlate his action with the state as effectively as the good type, which introduces an additional reputational cost to allowing.

The beliefs about the true underlying state will be strongly influenced by actions taken so that often, a player's reputation may suffer despite the fact that the public retrospectively judge his decision to be most likely correct. So there is often dissonance between beliefs about the quality of the agent, and the quality of his decisions. For example when the decision maker bans, the public will typically be persuaded that the activity is unsafe, but it still damages his reputation. Banning is interpreted as a sign of weakness, even though the public think he did right. This kind of dissonance does not arise in the boundedly rational model since by construction the DM is judged to be good precisely if his action appears to be the right one.

2.7 Conclusions

Understanding the effects of public accountability on the incentives of public officials has never been more relevant than it is today, with increasing demand for accountability in every area of public life. More positions in public office are being put up for election (e.g. police commissioners in the UK); politicians are reverting to public consultations on complex policy decisions such as the country's nuclear programme or membership of the EU; public employees are increasingly exposed to judgement by service users (e.g. malpractice litigation against physicians in the US). Greater accountability brings undeniable benefits in terms of better scrutiny of activities where there may be conflicts of interest, and more flexibility to oust under-performing officials. However, it also brings some costs and this paper is an attempt to understand the types of distortions that will arise and how the institutional design changes things.

I present a model for how risk regulators are held to account by the public and the effects of this accountability on policy choice. I propose a heuristic approach to accountability. This is motivated by substantive evidence from the field of psychology that individuals suffer from persistent biases and rely on simple rules in their perceptions of risk, and additionally on suggestive evidence that in complex environments individuals judging performance resort to shortcuts. The public will judge regulators by a simple retrospective assessment: how good do past policies seem today, in light of observed events?

I derive results for a number of dynamic problems, varying mainly the time horizon of the decision maker, and the characterisation of the solution is robust across all of them. The policy maker will attempt to match up public beliefs and his record of policy. This creates dynamic incentives to pander and to manipulate learning about risks, whether by preventing learning or encouraging it.

These incentives provide a rich picture of the paths of decision making, as the beliefs and policy record evolve. The analysis suggests, for example, that deregulation is more fragile in early days because the decision maker is more willing to change course in the case of a bad shock, as it is easier for him to change his record. Regulators with a strong history of supporting some activity or industry can have reputational reasons for being stubborn in their support after disaster happens, because it is harder for them to shake off their anti-regulation stance.

This may help to explain, for example, why different countries with similar information and circumstances have arrived at very different nuclear power policies. In Italy, the government was slow to get off the ground and by the 1980s only a handful of plants existed or were under construction; therefore, after the event at Chernobyl, the government did not find it too costly to reverse its stance on nuclear power: all existing reactors were decommissioned. The experience in France was very different: in 1974 Prime Minister Pierre Messmer announced a huge nuclear power programme aimed at generating all of France's electricity from nuclear power, and indeed France installed 56 reactors over the next 15 years. This programme continued largely undisturbed after Chernobyl: policy makers persevered and waited for public alarm to subside.

The model also suggests that we are more likely to find moderate policy approaches when the public support deregulation and more extreme ones (either for or against regulation) when public concern about the activity is high. Deregulation always involves reputational risk because beliefs may change. If the public support deregulation, the decision maker may prefer to maintain a moderate record[13] in case beliefs are reversed. Also, by introducing only limited deregulation the policy maker reduces the probability of an event that may change beliefs. On the other hand, if the public perceive the risk as high, the regulator will not benefit from intermediate levels of regulation. He will either try to persuade the public that the underlying risk is low, by deregulating and letting the public see that disaster is very rare, or he will try to quickly erase his association with the activity by banning the activity entirely.

My results also point to a strong bias towards the safe action (tight regulation). For any non-zero level of risk the decision maker will prefer to ban for certain initial states, and he will ban forever. Educating the public is not in itself of any value to a regulator in these circumstances. In the Bayesian model, risk-taking would be rewarded: a regulator that chooses to allow and is succesful would always have the highest reputation in equilibrium. Here, that is not the case; the public are perfectly content with a regulator that does not challenge their beliefs and regulators will be complacent in prohibiting activities that are feared by the public. This provides one

[13]We can also interpret the moderate record as persistently choosing intermediate levels of regulation rather than mixing between allowing and banning.

explanation, for example, for why some drugs are almost universally illegal despite many experts supporting legalisation. As long as the public believe that the risks of legalisation outweigh the benefits, they will praise the decision to ban and their beliefs will be preserved. Beliefs may change over time through other channels, but a regulator who cares only about his reputation, will not benefit from changing perceptions.

The model provides a framework to analyse other questions. I focus on understanding the effects of a number of institutional variations. The analysis suggests that increasing the amount of information that the public observe, by providing them with additional "outside" signals, can improve welfare. The largest effect will be a large reduction in conservatism, as the decision maker no longer has a monopoly over the learning of the public and banning ceases to preserve beliefs. Extending the length of terms for decision makers causes more conservatism when the public favour strict regulation; when they favour deregulation, longer office terms lead to better policies for the extremes of low and high risk. Finally, I consider the effect of limiting policy flexibility. Costly switching prevents decision makers resorting to banning after disasters; however, it increases ex ante conservatism and the welfare effects are ambiguous.

Chapter 3

Incumbency Effects and Ideological Extremism in Elections

I propose a model of incumbency effects where political candidates vary in quality and compete over policy platforms. Any kind of electoral bias will bias policies towards the favoured candidate. However, I find that an electoral bias against the incumbent leads to more extreme ideological positions than one in his favour, because of the effect on polarisation. When the incumbent is systematically favoured this makes the value of winning to parties higher, giving them a dynamic incentive to compete more fiercely and close the policy gap. Incumbency effects arise endogenously when (rational) voters observe imperfect information about candidates' competence. When the distribution of quality is skewed towards poor quality, two effects (selection and exposure) act in opposite directions. When the information about challengers is of too low a quality, selection may be very weak and incumbency disadvantage may arise.

3.1 Introduction

It has been observed that many democracies display incumbency effects, systematic electoral advantages (or disadvantages) for incumbents. Incumbency effects will affect how competitive elections are and therefore the policy choices of competing parties. For example, if candidates expect that winning hurts their chances in tomorrow's electoral race, and popularity tends to swing systematically between the parties, they will be less willing to compromise on their policy offering in order to attract votes; instead they will simply wait for the tide to turn their way.

I propose a model of political competition on a one-dimensional policy space between two policy-motivated parties, which have ideal policies on either side of the median's preferred policy. Policy platforms are perfectly observed and there is full commitment to them. Politicians are also characterised by their competence, which has value for voters but is not perfectly observed. Voters have some information about the competence of a new candidate; they may know about his credentials from past positions held and may observe him during the electoral campaign. For the incumbent, voters have some additional information from his track record in office. They vote for the candidate who gives them higher overall utility, considering both their competence and policy offering.

First, I analyse the effect that incumbency biases have on policy choice. Previous literature on the static probabilistic model of voting has shown that a bias towards one of the candidates will bias average policy in favour of the more popular party, and that uncertainty over today's electoral outcomes will increase policy polarisation (Lindbeck and Weibull 1993). Both these types of distortion will increase overall policy variance. The question I address is what happens in a dynamic setting, where elections are repeated and the bias depends on the outcomes of previous elections. I find that a bias against the incumbent is particularly detrimental in terms of generating larger policy variance.

Holding office has value, because it allows the governing party to set policy closer to its ideal than the rival party would. Conditional on winning, choosing a platform closer to the party's ideal gives the party higher utility. But by moving away from the median's ideal, the party hurts its chances of winning and it becomes more likely to end up with the less desirable policy of the rival party. This provides the usual trade-off in probabilistic voting models.

In the dynamic context, the value of increasing today's probability of winning depends on whether winning today gives the party an edge over future opponents. When there is an electoral bias that favours the incumbent, there is an added benefit to offering a more competitive platform: winning today means tomorrow one has to fight less hard to win office. When elections are biased against the incumbent, the incentive to compromise on policy is weakened, now there is a cost to winning today. Even though the net effect of any electoral bias is to increase policy variance, the dynamic effect described implies that variance increases more steeply with a bias in favour of the challenger than a bias in favour of the incumbent.

I next consider how incumbency effects could arise endogenously. The question of what drives incumbency effects is a particularly interesting question in the light of some recent empirical work which has found that developing countries often display incumbency disadvantage. Several models have been proposed to explain the strong incumbency advantage typical of public offices in Western countries, particularly the US. So can we find an explanation for both incumbency advantage and disadvantage? What would determine which democracies fall into each case?

One mechanism that has already been shown to generate endogenous incumbency effects is voter learning about an imperfectly observed dimension of candidate quality (Ashworth and Bueno de Mesquita 2008). I show that if we generalise the model and allow for skewed distributions of quality, then the structure of information, specifically differences in the quality of signals generated on the job versus outside the job, is crucial in determining whether an incumbency advantage or disadvantage emerges. When the distribution of qualities is skewed to the right (so that it is more likely for a candidate to have below average quality than above average), two effects will arise from voter learning that go in opposite directions.

The selection effect, standard in the literature on incumbency advantage, is a positive filtering of candidates - elected candidates reached office because they had a (weakly) better reputation than their opponent. Therefore in future elections they should perform better than a randomly picked new challenger. However, being in office has an exposure effect, since it makes more information about the candidate's quality available to voters. This could mean that, if one performs badly, voters would rather vote in a challenger that they know little about rather than the incumbent who they are more convinced is bad.

When it is more likely that a candidate has above average quality, both effects go towards creating an incumbency advantage: selection does the initial filtering, and exposure only helps give the incumbent more chances to display his (likely) higher quality. But when low quality candidates are more common than high ones, the exposure effect is negative: incumbents have more chances of being exposed as mediocre. In this case, an incumbency advantage emerges only if the selection filter is sufficiently good. This translates into a condition on the accuracy of the different types of signals: incumbency advantage emerges when the signals for new candidates are sufficiently accurate relative to those produced by an incumbent in office.

The empirical literature in political science measures the incumbency advantage: the difference in probability of winning between incumbent politicians running for re-election and candidates running in an open seat election. Note that earlier I referred to an incumbency *bias*, the electoral lead the incumbent has in a given election. In the endogenous model, this bias will not be fixed but will depend on a state (the track record of the incumbent). The incumbency advantage is defined as the conditional expectation of the incumbency bias (conditional on reaching office). In my combined model of competition with ideological policies and competence, I find that polarisation depends on the *unconditional* incumbency bias.

Section 3.2 presents a model where candidates and voters have ideological preferences on a one-dimensional policy space, and consider how a fixed incumbency bias affects competition over policies. Section 3.3 presents a model for endogenous incumbency effects based on differential learning about incumbents and challengers in an election where candidates vary in quality. Section 3.4 combines the two elements, competition over ideological policies and competence, and presents some insights about welfare. Section 3.5 concludes.

Related Literature

The rise of ideological polarisation in political competition has been gaining attention in the political science literature. McCarty, Poole, and Rosenthal 2006 consider rising polarisation in the US, and show a correlation with growing inequality. Feddersen and Gul 2014 propose a model in which election outcomes depend on support by voters and donors. Uncertainty over the weight of each of the groups generates polarisation. Prato 2011 and Nunnari and Zápal 2014 also propose dynamic mechanisms that generate ideological extremism.

Over the last 50 years, a large empirical literature has looked at the effect of incumbency on prospects for re-election. This literature initially focused entirely on developed countries, and mostly on the US House of Representatives[1]. The literature consistently found a strong incumbency advantage in elected offices.

Zaller 1998 suggests this might be driven by voter learning and Ashworth and Bueno de Mesquita 2008 formalise this intuition. They propose a model of selection and strategic entry, where voters observe information about candidates prior to an election, and elect candidates with better signals of quality. Therefore, once in office

[1] It is not possible to provide an exhaustive list of references, but some notable contributions include Alford and D. W. Brady 1989, Ansolabehere and Snyder 2002, Jewell and Breaux 1988, Cox and Katz 1996, Cox and Morgenstern 1993, Ferejohn 1977, Gelman and King 1990.

they are likely to outperform challengers who have not been through the electoral filter. Further, opposition parties may not find it worthwhile to search for new candidates, if the incumbent has a strong reputation and search is costly. These results are shown for the case in which the underlying distribution of types (for all candidates) is normal and therefore symmetric.

More recently, evidence has emerged in sharp contrast to earlier findings: it appears that the pattern of incumbency effect is precisely the opposite across much of the developing world, where incumbents suffer from an electoral disadvantage. Molina 2001 confirms this for much of Latin America, Linden 2004 and Uppal 2009 find the same result in state legislatures in India, and Titiunik 2009 confirms a similar picture for Brazil's mayoral elections.

In my model of learning about candidate quality, I borrow the underlying framework from Ashworth and Bueno de Mesquita 2008 with two modifications. First, I allow for the distribution of candidate quality to be skewed so that the pool of candidates may have more candidates with below (or above) average quality. Second, I allow for signal accuracy to potentially differ between incumbents and challengers. An incumbent must make decisions in office, and the outcomes of those decisions are observable to voters (even if there is substantial noise in the outcomes). Section 3.3 demonstrates how, with this simple departure from Ashworth and Bueno de Mesquita 2008, one can predict both incumbency advantage and disadvantage in a system in which imperfectly informed Bayesian voters elect office holders by majority rule.

Dewan and Hortala-Vallve 2013 consider how learning about the competence of incumbents and challengers may affect the incentives for candidates in office to carry out risky reforms. Their environment also has the feature that the information technology for learning about incumbents may be different to that of challengers. During an electoral campaign, voters are able to learn about some of the candidates' skills, for example, whether he is a good communicator[2] or has relevant knowledge for the job [3]. However, there is also information that voters can only observe about a candidate once he is in the job, for example, how authoritatively he deals with a crisis[4], or how successful he is at delivering his electoral promises[5].

[2]e.g. Nicola Sturgeon, leader of the Scottish National Party, enjoyed a significant boost to her approval ratings after an impressive performance at the leaders' debate in the run-up to the 2015 UK election.

[3]e.g. Natalie Bennett, leader of the UK's Green Party, suffered strong criticism when, during a radio interview, she failed to know the details relating to one of her party's main electoral pledges.

[4]e.g. As the UK has come out of recession, the ruling Conservative party's lead over Labour on the question of "most trusted team with the economy" reached 20% in late 2014

[5]e.g. Nick Clegg, leader of the Liberal Democrats, suffered heavy losses in voter support, after failing to prevent increases in university tuition fees while in coalition with the Conservatives.

3.2 Ideological Polarisation in the Presence of Incumbency Effects

3.2.1 Competition over Ideological Policies

I consider a one-dimensional policy space $x \in (-\infty, \infty)$. A continuum of voters have single-peaked preferences over policies; voter i's ideal policy is denoted $\widehat{x}^i$. I assume a linear-loss utility function: $-\alpha^v \mid x - \widehat{x}^i \mid$. The distribution of voters' ideal policies can have any (continuous) distribution $F(\widehat{x}^i)$. For now the only relevant parameter of the distribution will be the median bliss point, denoted $\widehat{x}^m$.

Two candidates, $J \in \{L, R\}$, compete for office. Their ideal points lie on opposite sides of the median bliss point: $\widehat{x}^L < \widehat{x}^m < \widehat{x}^R$. The candidates also have linear single-peaked utility: $-\alpha^C \mid x - \widehat{x}^J \mid$. I allow for the marginal utility to be different between voters and candidates but assume it is the same for the two candidates. Further, I assume that candidates care about the policies implemented, regardless of the party implementing them. I assume there are no office rents.

For any given pair of proposed platforms (x_L, x_R), there is some uncertainty over the identity of the winner. I model this as a stochastic aggregate taste shock[6]. I assume that the distribution of this taste shock depends only on whether the candidate is an incumbent or a challenger.

Voters are assumed to be fully myopic. The utility voter i obtains from candidate L winning the election is the policy payoff from L's policy plus some shock which reflects public taste about orthogonal dimensions of quality such as competence, charisma, etc: $-\alpha^v \mid x_L - \widehat{x}^i \mid + \varepsilon$. The utility obtained from R winning is $-\alpha^v \mid x_R - \widehat{x}^i \mid$.

Let me define the state of the world as the identity of the incumbent in the election taking place at the beginning of the period: $\tilde{J} \in \{L, R\}$. For simplicity, I assume $\varepsilon \sim U\left[b - \frac{1}{2\phi}, b + \frac{1}{2\phi}\right]$, where b is an incumbency-related bias (i.e. a function of the incumbency state). Formally,

$$b = \widehat{b} I \left\{ \tilde{J} = L \right\} - \widehat{b} I \left\{ \tilde{J} = R \right\}$$

When $\hat{b} > 0$, the distribution of the taste shock ε has a positive mean whenever L is incumbent, giving him a lead in the electoral race, and the shock has a negative mean when he is the challenger (so again there is a bias towards the incumbent).

[6] Other interpretations are possible, such as candidates being uncertain about the exact position of the median voter (Lindbeck and Weibull 1993).

When $\hat{b} < 0$, we have the converse case, the electoral bias works systematically against the incumbent.

The electoral rule is simple majority with ties resolved by a fair coin toss. The game proceeds as follows. At the beginning of a period, given the identity of the incumbent, candidates simultaneously choose their platforms which are fully binding. Voters observe the platform offers and vote sincerely. A majority vote winner is declared, whose policy is implemented and who becomes the incumbent for the next period.

I analyse a static model first, then proceed to the the dynamic model, where elections are infinitely repeated. In the dynamic environment, I assume voters are myopic and candidates are forward-looking and discount future payoffs at a rate $\beta \in [0, 1]$.

3.2.2 Results in a Static Model

Since there are no idiosyncratic party-preference shocks, the ordering of voters' preferences over the two parties (given their platforms) is always the same, irrespective of the aggregate taste shock. This implies that the median voter's preference over the parties determines the winner of the election since for any ε:

If $\alpha^v \mid x_L - \widehat{x}^m \mid > -\alpha^v \mid x_R - \widehat{x}^m \mid + \varepsilon$ then $-\alpha^v \mid x_L - \widehat{x}^i \mid > -\alpha^v \mid x_R - \widehat{x}^i \mid + \varepsilon$
for all i such that $\widehat{x}^i < \widehat{x}^m$, and
if $-\alpha^v \mid x_L - \widehat{x}^m \mid < -\alpha^v \mid x_R - \widehat{x}^m \mid + \varepsilon$ then $-\alpha^v \mid x_L - \widehat{x}^i \mid < -\alpha^v \mid x_R - \widehat{x}^i \mid + \varepsilon$
for all i such that $\widehat{x}^i > \widehat{x}^m$.

Lemma 8. *Candidates will never choose policies on the opposite side of the median bliss point to their own bliss point:* $x_L \leq \widehat{x}^m \leq x_R$.

This is trivial to show. Suppose $x_L > \widehat{x}^m$. Then there exists another policy $x'_L < \widehat{x}^m$ that is preferred to R's policy by at least as many voters as when he chose x_L, and is strictly preferred by L himself.

Lemma 9. *Candidates will not choose policies that are more extreme than their own bliss points:* $x_L \geq \widehat{x}^L$ *and* $x_R \leq \widehat{x}^R$.

The proof is very similar to the proof of Lemma 8. There exists a deviation on the other side of his own bliss point that is preferred to R's policy by strictly more voters and again gives the candidate higher utility conditional on winning.

The probability of L winning, given platforms (x_L, x_R) is

$$\begin{aligned} p_L(x_L, x_R) &= \Pr\left[-\alpha^v\left(\widehat{x}^m - x_L\right) + \varepsilon > -\alpha^v\left(x_R - \widehat{x}^m\right)\right] \\ &= \Pr\left[\varepsilon > \alpha^v\left(2\widehat{x}^m - x_L - x_R\right)\right] \\ &= \Pr\left[\varepsilon > 2\alpha^v\left(\widehat{x}^m - \overline{x}\right)\right] \\ &= \phi\left[2\alpha^v\left(\overline{x} - \widehat{x}^m\right) + b + \frac{1}{2\phi}\right] \end{aligned} \tag{3.1}$$

where $\overline{x} = \frac{x_L + x_R}{2}$ is the mean policy offered. From this expression we see that the probability of L winning is increasing in both x_L (as his policy approaches that of the median) and in x_R (as the opponent's policy moves further from the median). We can also see that when $\hat{b} > 0$, the bias favours the incumbent; e.g. when L is the incumbent $b = \widehat{b} > 0$ and so this term increases p_L. When $\hat{b} < 0$, L's chances of winning take a hit when he is in office, i.e. the bias favours the challenger.

Using lemmas 8 and 9, L's problem is

$$\max_{x_L \in [\widehat{x}^L, \widehat{x}^m]} -p_L\left(x_L, x_R\right)\alpha^C\left(x_L - \widehat{x}^L\right) - \left(1 - p_L\left(x_L, x_R\right)\right)\alpha^C\left(x_R - \widehat{x}^L\right) \tag{3.2}$$

R's problem is

$$\max_{x_R \in [\widehat{x}^m, \widehat{x}^R]} -p_L\left(x_L, x_R\right)\alpha^C\left(\widehat{x}^R - x_L\right) - \left(1 - p_L\left(x_L, x_R\right)\right)\alpha^C\left(\widehat{x}_R - x_R\right) \tag{3.3}$$

The first order condition for L is:

$$\alpha^C\phi\left[2\alpha^v\left(\overline{x} - \widehat{x}^m\right) + b + \frac{1}{2\phi}\right] = \phi\alpha^v\alpha^C\left(x_R - x_L\right) \tag{3.4}$$

The left hand side of this is the marginal cost to L of moving marginally closer to the median since, conditional on winning, he obtains a lower utility. The right hand side is the marginal benefit, the increase in his probability of winning times the utility difference between his and his opponent's policies. The first order conditions for R has a similar form:

$$\alpha^C\phi\left[-2\alpha^v\left(\overline{x} - \widehat{x}^m\right) - b + \frac{1}{2\phi}\right] = \phi\alpha^v\alpha^C\left(x_R - x_L\right) \tag{3.5}$$

Equations (3.4) and (3.5) implicitly define the best response functions of L and R. Combining the two conditions yields

$$\overline{x} = \widehat{x}^m - \frac{b}{2\alpha^v} \tag{3.6}$$

The bias in the popularity shocks leads to a bias in the average policy offered (and implemented), where this bias is increasing in the magnitude of the popularity shock and decreasing in the intensity of policy preferences. This is intuitive: these two parameters determine the relative importance of policy and non-policy considerations. When voters care a lot about policy (high α^v) and the bias is smaller, the advantaged party will not be able to exploit its advantage as much and policies will be closer to the median's ideal.

I solve for the difference between the platforms of the two parties in equilibrium ($\delta x = x_R - x_L$), which I will call the ideological polarisation from now on.

$$\delta x = \frac{1}{2\phi\alpha^v} \tag{3.7}$$

The degree of polarisation increases with ϕ because of the usual argument: high ϕ means that the size of popularity shocks is small and hence elections are more often determined by platforms. Therefore, candidates' incentive to compromise on policy in order to attract votes is stronger. The bias in popularity shocks does not affect polarisation in the one-shot game.

It is straightforward to find the actual equilibrium strategies (x_L, x_R) from the system for $\overline{x}$ and δx. Overall policy variance depends on both how far average policy is from the median, and how polarised the parties' platforms are around the average policy. I define policy variance as the distance from the median of the furthest policy platform, $Var(x) = \max|\hat{x}^m - x_J|$. Then,

$$Var(x) = \frac{|\hat{b}|}{2\alpha^v} + \frac{1}{4\phi\alpha^v} \tag{3.8}$$

3.2.3 Results in a Dynamic model

For simplicity, I consider a perfectly symmetric model. Specifically, I assume $\widehat{x}^m = 0$, and $\widehat{x}^m - \widehat{x}^L = \widehat{x}^R - \widehat{x}^m$.

The problem has a stationary structure. I look for stationary Markov perfect equilibria, where strategies are only a function of the state. Therefore, the solution is characterised by two pairs of policies: $\left\{\left(x_L^L, x_L^R\right), \left(x_R^L, x_R^R\right)\right\}$ where $x_J^{\tilde{J}}$ is the policy chosen by candidate J when party $\tilde{J}$ is the incumbent.

The Lemmas presented in the previous section hold true in the repeated environment. Parties gain nothing from overshooting their policy beyond the median's bliss point or beyond their own.

Denote by $V_J\left(\tilde{J}\right)$ the value function of candidate J when, entering the election for that period, $\tilde{J}$ is the incumbent. One can write the candidates' problem recursively. L must solve

$$\begin{aligned}\max_{x_L \in [\widehat{x}^L, 0]} & - p_L\left(x_L, x_R^{\tilde{J}}, \tilde{J}\right) \alpha^C \left(x_L - \widehat{x}^L\right) - \left(1 - p_L\left(x_L, x_R^{\tilde{J}}, \tilde{J}\right)\right) \alpha^C \left(x_R^{\tilde{J}} - \widehat{x}^L\right) \\ & + \beta \left[p_L\left(x_L, x_R^{\tilde{J}}, \tilde{J}\right) V_L(L) + \left(1 - p_L\left(x_L, x_R^{\tilde{J}}, \tilde{J}\right)\right) V_L(R)\right]\end{aligned} \tag{3.9}$$

R must solve,

$$\begin{aligned}\max_{x_R \in [0, \widehat{x}^R]} & - p_L\left(x_L^{\tilde{J}}, x_R, \tilde{J}\right) \alpha^C \left(\widehat{x}^R - x_L^{\tilde{J}}\right) - \left(1 - p_L\left(x_L^{\tilde{J}}, x_R, \tilde{J}\right)\right) \alpha^C \left(\widehat{x}^R - x_R\right) \\ & + \beta \left[p_L\left(x_L^{\tilde{J}}, x_R, \tilde{J}\right) V_R(L) + \left(1 - p_L\left(x_L^{\tilde{J}}, x_R, \tilde{J}\right)\right) V_R(R)\right]\end{aligned} \tag{3.10}$$

I find that polarisation is larger and overall policy variance rises faster, when there is a systematic electoral bias towards challengers, than when the bias favours the incumbent.

Proposition 7. *In the dynamic model where the taste shock has a bias $\hat{b}$ towards the incumbent, mean policies are $\overline{x}^L = -\frac{\widehat{b}}{2\alpha^v}$ and $\overline{x}^R = \frac{\widehat{b}}{2\alpha^v}$; polarisation is $\delta x = \frac{1}{2\phi\alpha^v} - \frac{\beta\hat{b}}{\alpha^v}$. Overall policy variance rises at a rate of $\frac{1-\beta}{2\alpha^v}$ with respect to $\hat{b}$ when $\hat{b} > 0$, and at a rate of $\frac{1+\beta}{2\alpha^v}$ with respect to $|\hat{b}|$ when $\hat{b} < 0$.*

Proof in the Appendix.

In the dynamic setting, there are two components to the degree of ideological polarisation. There is the static term from today's uncertainty over voters' preferences (the taste shock) that leads to weaker competition. The new term depends on the expected marginal continuation value of winning. In this case, who wins today's election determines who enjoys an electoral advantage tomorrow. In equilibrium, $p_L = \frac{1}{2}$ in every period. So, the effect winning today won't affect the probability of winning beyond the immediate next period. It will, however, affect tomorrow's advantage and how hard one has to compete in that period's election. This translates into a bias in the average policy tomorrow. The marginal continuation value of winning is therefore the (discounted) difference in mean policy between the two states (winning and losing): $\frac{\beta\hat{b}}{\alpha^v}$.

Consider first the case in which the incumbent enjoys an electoral advantage, $\hat{b} > 0$. In this case, there is a dampening of the static effect on polarisation from dynamic considerations. Today's uncertainty over the popularity shock makes the parties compete less hard for votes because, to some extent, the winner may be determined by factors out of their control. But parties also consider the future: if they lose today, they will continue to be disadvantaged tomorrow; if they win, they will be ahead in the race tomorrow. This additional cost of losing means that the parties will fight a bit harder to attract votes by moving their platform closer to the median's ideal policy.

When the election has a bias toward the challenger, $\hat{b} < 0$, the dynamic effect widens the gap between the parties' platforms. Now losing an election has a positive side to it - the continuation value is higher for the challenger than the incumbent because tomorrow the loser has the better electoral prospects. Losing today has a guarantee that the future is brighter which further reduces competitive pressures.

Overall variance in policy always rises (for $\beta < 1$) when the magnitude of the electoral bias goes up. A larger electoral bias primarily increases the distance of mean policy from the median. Anticipation of how the future mean policy will be biased then has an indirect effect on parties' willingness to compromise today.

Therefore, when a political system is characterised by a large electoral bias towards the incumbent, the model predicts that one should observe large swings in the centre ground ($\bar{x}$) as parties switch over, but both parties should be rather closely positioned ideologically (small δx). In contrast, a bias against the incumbent should generate large swings in the centre but also a more pronounced divide between the parties' positions at every election.

The discount factor has no effect on the average bias of the policies. Instead the degree of polarization δx is strictly decreasing (increasing) in β when there is a bias towards (against) the incumbent. As β grows, candidates become more patient. If they expect the incumbent to be favoured tomorrow this means higher β increases the marginal value of winning and so competition will be stronger, reducing the wedge between the platforms (and vice versa).

3.3 Endogenous Incumbency Effects

The empirical literature on incumbency effects has identified both democracies with an incumbency advantage and with incumbency disadvantage. I modify the framework in Ashworth and Bueno de Mesquita 2008 and demonstrate that with a skewed distribution of candidate quality types, incumbency disadvantage as well as incumbency advantage may arise endogenously. I demonstrate that differences in observed incumbency effects could be explained by differences in information technology. I simplify the type space and assume quality is binary, as are signals. This keeps the updating of beliefs analytically tractable.

3.3.1 A model of learning about candidate quality

There are two parties, $J \in \{L, R\}$. In each period, $t \in \{1, 2\}$, each party puts forward a candidate. Candidates differ in quality, $v_J \in \{0, 1\}$, the quality of a candidate is constant over time and quality cannot be observed directly by voters. The prior on quality is $Pr\left(v_J = 1\right) = \mu$ for candidates from either party. This implies $E\left(v_J\right) = \mu$, and $Pr(v_J < \mu) > \frac{1}{2}$ if and only if $\mu < \frac{1}{2}$. That is, if $\mu < \frac{1}{2}$, more candidates are bad (below average quality) than good (above average quality).

Voters observe signals of candidates' quality, and these signals may have different accuracy when the candidate is in office than when he is an inexperienced challenger. Let the signal observed in period t for candidate J be denoted $s_t^J \in \{0, 1\}$. For the incumbent, with probability p, the true quality is revealed and with probability $1 - p$ voters get a signal that is uncorrelated with v_J and is an independent draw from the same distribution as that of v_J. Instead for the challenger, with probability q the signal is perfectly correlated with the type, and again otherwise the signal is independently and identically distributed as v_J.

Then for the incumbent's signal:

$$\begin{array}{ll} \Pr\left(s_t^J = 1 \mid v_J = 1\right) = p + (1-p)\,\mu & \Pr\left(s_t^J = 0 \mid v_J = 1\right) = (1-p)\,(1-\mu) \\ \Pr\left(s_t^J = 1 \mid v_J = 0\right) = (1-p)\,\mu & \Pr\left(s_t^J = 0 \mid v_J = 0\right) = p + (1-p)\,(1-\mu) \end{array}$$

Signal probabilities for the challenger are of the same form, but with accuracy q.

These assumptions about the signals imply that the unconditional distribution of signals observed is the same as the distribution of underlying qualities, but the two variables are not perfectly correlated so that signals are imperfectly informative. Therefore p, q indicate the degree of correlation of signals with actual quality, and higher p, q correspond to more informative signals.

Quality is valuable to voters and they will select the candidate with highest expected quality. I will assume that when a party loses the election, it automatically and costlessly picks a new candidate from its pool.

In the first period, both candidates are challengers, and the public observe a signal about each candidate. Given the signals, they form an assessment about the expected quality of the candidate. They vote for the candidate with highest expected quality, and ties are broken by a fair coin toss. In the second period, voters observe a new signal for the incumbent and also a signal for the new challenger, and again vote for the candidate with highest expected quality, based on all signals observed up to that point. Let the expected quality of a candidate be denoted $\tilde{\mu}_t^J$, which will be a function of all the signals available for candidates in the running.

Definition (Incumbency Advantage) The incumbency advantage is the expected difference between the probability of a party winning the election with an incumbent and the probability of a party winning in an open seat election.

I will denote the incumbency advantage by Δ, so we have

$$\Delta = Pr\,(win \mid incumbent) - Pr\,(win \mid open\ seat)$$

Since the priors are the same for the two parties, $Pr\,(win \mid open\ seat) = \frac{1}{2}$.

3.3.2 Two effects: selection and exposure

The information structure I have assumed gives rise to two very natural effects for the electoral prospects of a party. In this section, I show each effect in isolation by focusing on particular parametric assumptions.

The pure selection effect: $p = 0$

Suppose that the only information voters ever have about a candidate is the signal he produced when he first stood for the open seat election, and that once in office they learn nothing about his ability.

Lemma 10. *The pure selection effect is always positive: when $p = 0$ and $q > 0$, we have $\Delta > 0$ for all μ, q.*

Proof. The probability of winning is $\frac{1}{2}$ for each party in the open seat election. Assume without loss of generality that L wins the first election.

Now consider the second period. Since $p = 0$ the second signal produced by L is entirely uninformative and the posterior remains unchanged from what it was when the first election took place. In addition, an informative signal will be observed for the new challenger R. Since this signal is as accurate as that from $t = 1$ for the incumbent, the voters will re-elect L if $s_1^L > s_2^R$, they elect R if $s_1^L < s_2^R$, and toss a coin if $s_1^L = s_2^R$.

$$\begin{aligned} &Pr\,(L \text{ wins at } t = 2 \mid L \text{ incumbent}) \\ &\quad = Pr\left(s_1^L = 1 \mid L \text{ wins at } t = 1\right) \cdot Pr\left(L \text{ wins at } t = 2 \mid s_1^L = 1\right) \\ &\quad + Pr\left(s_1^L = 0 \mid L \text{ wins at } t = 1\right) \cdot Pr\left(L \text{ wins at } t = 2 \mid s_1^L = 0\right) \end{aligned} \tag{3.11}$$

Regardless of q, the probability of $s_2^R = 1$ is μ. This implies

$$Pr\left(L \text{ wins at } t = 2 \mid s_1^L = 1\right) = \left[\frac{\mu}{2} + 1 - \mu\right]$$

$$\text{and } Pr\left(L \text{ wins at } t = 2 \mid s_1^L = 0\right) = \left[\frac{1-\mu}{2}\right]$$

By Bayes' rule,

$$Pr\left(s_1^L = 1 \mid L \text{ wins at } t = 1\right) == \frac{\mu\left(\frac{\mu}{2} + 1 - \mu\right)}{\frac{1}{2}}$$

$$\text{and } Pr\left(s_1^L = 0 \mid L \text{ wins at } t = 1\right) = (1-\mu)^2$$

Note that for all μ, $Pr\left(s_{L,1} = 1 \mid L \text{ wins at } t = 1\right) > \mu$. Since candidates with better signals are retained, incumbents on average have a better record than can be generated by a random candidate. Finally,

$$Pr\left(L \text{ wins at } t = 2 \mid L \text{ incumbent}\right) = 2\mu\left(1 - \frac{\mu}{2}\right)^2 + (1-\mu)^2\left(\frac{1-\mu}{2}\right)$$

and the incumbency effect is

$$\Delta = \frac{\mu(1-\mu)}{2} > 0$$

□

This is the selection effect - elections provide a filter to let higher quality candidates into office. Selection into office is more likely with high signals, and is therefore related to better electoral prospects at future contests. Random challengers should on average produce worse signals than a candidate that has already passed the test previously, generating an incumbency advantage.

The pure exposure effect: $q = 0$

Consider the other extreme next. Suppose that candidates can only demonstrate their ability by doing the job, and challengers' signals are entirely uninformative.

Lemma 11. *The pure exposure effect is positive if and only if median quality is higher than average quality: when $q = 0$, we have $\Delta \geq 0$ if and only if $\mu \geq \frac{1}{2}$, for all $p > 0$.*

Proof. In period 1, voters have no information and so there is a tie with probability

1 and each candidate wins with probability $\frac{1}{2}$. Assume without loss of generality that L wins the election in period 1.

In period 2, again, voters can learn nothing about the challenger R, and their posterior belief is $\tilde{\mu}_2^R\left(s_2^R\right) = \mu$. However, they now get to see L on the job and this will reveal some useful information about his ability. By Bayes' rule, and for $p > 0$,

$$\tilde{\mu}_2^L\left(s_1^L, s_2^L = 1\right) = \frac{\mu\left(p + (1-p)\,\mu\right)}{\mu} > \mu > \frac{\mu\,(1-p)\,(1-\mu)}{1-\mu} = \tilde{\mu}_2^L\left(s_1^L, s_2^L = 0\right)$$

The incumbent is preferred to the challenger if and only if he produces a high signal. The probability of an incumbent winning re-election is precisely μ and we have

$$\Delta = \mu - \frac{1}{2} \geq 0 \text{ if and only if } \mu \geq \frac{1}{2}$$

□

This is the exposure effect: when more information is observed about incumbents than challengers, being in office can be a blessing or a curse. Which of these is true depends on whether it is more likely that candidates will be exposed as mediocre or as outstanding leaders. If candidates are more likely to have low quality, for example, then being in office is likely to hurt the candidate's chances, as the public remain agnostic about the challenger but they observe evidence of the incumbent's incompetence.

3.3.3 Incumbency: advantage or disadvantage?

How does the incumbency effect depend on the general information structure?

Remark. If the median quality of candidates is above the mean, $\mu \geq \frac{1}{2}$, there cannot be an incumbency disadvantage.

This is immediate given Lemmas 10 and 11: the selection effect is always positive , and the exposure effect is also positive if candidates are more likely to be good than bad.

Proposition 8. *Whenever the median quality is lower than the average quality, there is an incumbency disadvantage when the selection effect is "weak enough". Formally, for all $\mu < \frac{1}{2}$, and and $p \in (0,1)$, there exists a set of $q \in [0,1]$ with positive measure such that $\Delta < 0$ and includes, at least, an interval of the form $[0, q^*]$. The solution, in the case of binary quality and signals, has the following form. There exist values $\underline{\mu}$ and $\bar{\mu}$, with $0 < \underline{\mu} < \bar{\mu} < \frac{1}{2}$, such that:*

- *For $\mu < \underline{\mu}$, $p > 0$, we have $\Delta < 0$ for all $q < 1$.*
- *For $\mu > \bar{\mu}$, there exists a threshold $q^* \in (0, 1)$ such that $\Delta < 0$ for $q < q^* (\mu, p)$.*
- *For $\underline{\mu} < \mu < \bar{\mu}$ there are two thresholds q^* and q^{**}, with $q^* < q^{**}$, such that $\Delta \leq 0$ for $[0, q^*] \cap [q^{**}, 1)$.*

Proof in the Appendix.

If good leaders are rare, there is a tension between the selection and exposure effects. When selection is strong, i.e. when voters can effectively pick out good candidates in the open seat election, we should expect that incumbents are highly likely to be re-elected. Selection puts a filter on the candidates that get into office initially. Once in office, the exposure effect is acting on a population of incumbents with higher quality so that incumbents are likely to perform well. In contrast, when selection is weak, the filtering doesn't work effectively. A lot of mediocre candidates get a shot at the job, and most often are exposed to be bad leaders, losing at the second election.

The proposition distinguishes three cases, depending on the quality distribution (μ). The first case in Proposition 8 is where a very high proportion of candidates are bad. In this case the exposure effect is very strong, the electoral filter cannot sufficiently improve the average quality of candidates. In the first (open seat) election, the most likely outcome is to observe a bad signal from both candidates and let one through at random. Selection is poor despite potentially high q because with a single election most winners win because of ties, not because they performed well. In the second election, it is again quite likely that both the incumbent and the challenger produce bad signals. At this stage voters, though they are quite sure neither of them is any good, will opt for the challenger, simply because at least they have less evidence of his incompetence.

When the prior is below $\frac{1}{2}$, but not too low, then we have the most straightforward result. Above some level of precision for the signals of inexperienced candidates, selection is good enough to overcome exposure and generate incumbency advantage. When those signals are poor, so is selection, and incumbency leads to electoral disadvantage. Note that in the case of low q, it is perfectly rational for voters to frequently replace their leaders: good leaders are hard to come by and, further, can only be identified if they are given the chance to govern. It is an optimal strategy for voters to sample different candidates in search of the real stars.

The final case has incumbency advantage for intermediate values of q. The reason this case occurs is that incumbency advantage is not monotonic with respect to q because q influences both selection and exposure. Specifically, q affects the ranking of posteriors of incumbent and challenger in period $t = 2$ as well as the degree of correlation between the incumbent's first and second period signals. Higher q raises the probability of good signals of an incumbent, but it also makes the posteriors of the challenger more extreme. At q^* the ranking of posteriors $\tilde{\mu}(1,0)$ and $\tilde{\mu}(0)$ switches, so that there are more signal profiles for which the incumbent beats the challenger; but at q^{**} the posteriors $\tilde{\mu}(0,1)$ and $\tilde{\mu}(1)$ change order, so that now there are more cases where the challenger beats the incumbent[7].

3.4 Competition with Ideological Policies and Competence

Next, I combine the elements I have analysed independently in the previous two sections: ideological policies and candidate competence. This will enable me to show more clearly the difference between two related but distinct concepts - the expected marginal continuation value of winning and the incumbency advantage. In addition, I will comment on some welfare implications of the model.

I assume that some fraction of the electorate votes according to ideological considerations, and the rest vote based on which candidate has higher (expected) competence given observed signals for the two candidates running[8]. Candidates select their platforms for that period before the signals are realised.

Let a fraction $\gamma \in (0,1)$ of voters be ideological, and a fraction $(1-\gamma)$ be competence-based voters. All voters are myopic and vote sincerely. Candidate J wins the election if he gets a majority of the overall votes, with the size of the population normalised to 1.

Ideological voters have ideal points on the policy spectrum distributed according to $\hat{x}^i \sim U[-1,1]$. There is an aggregate taste shock that gives voters additional utility from L holding office, $\epsilon \sim U\left[-\frac{1}{2\phi}, \frac{1}{2\phi}\right]$. The indifferent voter satisfies

$$-\alpha^v(\hat{x}^{ind} - x_L) + \epsilon = -\alpha^v(x_R - \hat{x}^{ind})$$

and is therefore located at $\hat{x}^{ind} = \bar{x} + \epsilon/2\alpha^v$, for a given realisation of the shock. This would give L a fraction of the ideological vote equal to $F(\hat{x}^{ind}) = (\hat{x}^{ind}+1)/2$.

The election in the first period is an open seat election, and in the second

[7]See Figure 3.1a for an example of how the posteriors, and their ranking, change with q.

[8]Other interpretations are equivalent in terms of modelling assumptions, for example, that all voters value competence and whichever candidate has the better reputation experiences a shift (of fixed size) in the aggregate taste shock in his favour.

period one candidate will be an incumbent, while the other will be a challenger. Competence again is binary, and the technology for learning is the same as in Section 3.3, with the accuracy of signals for incumbents and challengers being potentially different. The competence-based voters will be able to observe all available signals before the election, and they will rationally update their beliefs about the quality of the candidates. The candidate with the higher reputation will win the competence-based vote, and the (competence-based) vote will be equally split in the case of a draw.

Period 2

At the start of period 2 there is an incumbent who carries forward an earlier signal from period 1. Suppose the incumbent in period 2 is L. Each candidate will choose a platform to maximise his expected policy utility, where expectations are now over both the taste shock and the new signals, which will affect who gets the competence-based votes.

For a given pair of platforms (x_L, x_R), the probability that L will win is[9]

$$p_L = \frac{1}{2} + 2\phi\alpha^v \bar{x}_2 + 2\phi\alpha^v \frac{(1-\gamma)}{\gamma} \left(Pr[\tilde{\mu}^L(s_1^L, s_2^L) > \tilde{\mu}^R(s_2^R)] - Pr[\tilde{\mu}^L(s_1^L, s_2^L) < \tilde{\mu}^R(s_2^R)] \right) \tag{3.13}$$

The probability that L will have the higher reputation in period 2 is conditional on s_1^L, which is known at the time that candidates choose their platforms for period 2. Let me use the following to economise on notation:

$$Pr[\tilde{\mu}^L \geq \tilde{\mu}^R | s_1^L] = Pr[\tilde{\mu}^L(s_1^L, s_2^L) > \tilde{\mu}^R(s_2^R)]$$

[9]The probability that L wins:

$$\begin{aligned}
p_L = E\Bigg\{ & Pr\left[\gamma\left(\frac{\bar{x}_2}{2} + \frac{\epsilon}{4\alpha^v} + \frac{1}{2}\right) + (1-\gamma)\mathbb{1}\left(\tilde{\mu}^L(s_1^L, s_2^L) > \tilde{\mu}^R(s_2^R)\right) + \frac{(1-\gamma)}{2}\mathbb{1}\left(\tilde{\mu}^L(s_1^L, s_2^L) = \tilde{\mu}^R(s_2^R)\right) \leq \frac{1}{2}\right] \Bigg\} \\
&= Pr[\tilde{\mu}^L(s_1^L, s_2^L) > \tilde{\mu}^R(s_2^R)] Pr\left[\gamma\left(\frac{\bar{x}_2}{2} + \frac{\epsilon}{4\alpha^v} + \frac{1}{2}\right) + (1-\gamma) \geq \frac{1}{2}\right] \\
&+ Pr[\tilde{\mu}^L(s_1^L, s_2^L) = \tilde{\mu}^R(s_2^R)] Pr\left[\gamma\left(\frac{\bar{x}_2}{2} + \frac{\epsilon}{4\alpha^v} + \frac{1}{2}\right) + \frac{(1-\gamma)}{2} \geq \frac{1}{2}\right] \\
&+ Pr[\tilde{\mu}^L(s_1^L, s_2^L) < \tilde{\mu}^R(s_2^R)] Pr\left[\gamma\left(\frac{\bar{x}_2}{2} + \frac{\epsilon}{4\alpha^v} + \frac{1}{2}\right) \geq \frac{1}{2}\right] \\
&= \frac{1}{2} + 2\phi\alpha^v \bar{x}_2 + 2\phi\alpha^v \frac{(1-\gamma)}{\gamma} \left(Pr[\tilde{\mu}^L(s_1^L, s_2^L) > \tilde{\mu}^R(s_2^R)] - Pr[\tilde{\mu}^L(s_1^L, s_2^L) < \tilde{\mu}^R(s_2^R)] \right)
\end{aligned} \tag{3.12}$$

L solves the problem

$$\max_{x_L} -\alpha^v(x_R - \hat{x}^L) + p_L(x_L, x_R, s_1)\alpha^v(x_R - x_L) \tag{3.14}$$

R solves the problem

$$\max_{x_R} -\alpha^v(x^R - \hat{x}_R) - p_L(x_L, x_R, s_1)\alpha^v(x_R - x_L) \tag{3.15}$$

The equilibrium policies will be such that

$$\bar{x}_2^L(s_1^L) = \frac{(1-\gamma)}{\gamma}\left(1 - 2Pr[\tilde{\mu}^L \geq \tilde{\mu}^R | s_1^L]\right) \text{ and } \delta x_2 = \frac{1}{2\phi\alpha^v} \tag{3.16}$$

When R is the incumbent, the degree of polarisation will be the same, and the mean policy will be $\bar{x}_2^R(s_1^R) = \frac{(1-\gamma)}{\gamma}\left(2Pr[\tilde{\mu}^L \geq \tilde{\mu}^R | 1 - s_1^R]\right)$.

In the last period, there are no dynamic considerations, so the polarisation term reduces back to the static effect from electoral uncertainty. The bias in policy will be equal to the expected electoral advantage in that period, which comes from winning the competence-based vote. With probability $Pr[\tilde{\mu}^L \geq \tilde{\mu}^R | s_1]$, L will win the competence-based vote, and with the complementary probability he will lose it.

If the incumbent carries forward a good signal, then he should continue to perform well today and with a high probability win the competence-based vote. Then, the average policy moves towards L's ideal policy which is negative. A bad signal brought forward will mean, in contrast, that he has a low chance of having the better reputation of the two and so average policy will be biased away from his ideal. Importantly, the incumbency bias is not fixed, it depends on the signal produced in period 1.

The probability that L will attract the competence-based vote, conditional on his first period signal, depends on the prior and accuracy of signals. As discussed in the proof for Proposition 8, the ranking of posteriors switches as p and q change, and this creates different cases for calculating the probability $Pr[\tilde{\mu}^L \geq \tilde{\mu}^R | s_1]$. For example, when q is very low and challenger signals are essentially uninformative, then L will win the competence-based vote if and only if he produces a good signal as incumbent.

Figure (3.1) provides an example, fixing the following parameter values: $\mu = 0.35$, $p = 0.5$, and $\gamma = 0.8$. Subfigure (3.1a) shows how the different posteriors are ranked when the accuracy of challenger signals improves. We can see that there will be two values of q at which the ranking of a posterior for L and a posterior for R switch. Subfigure (3.1b) shows the average policy platforms in the second

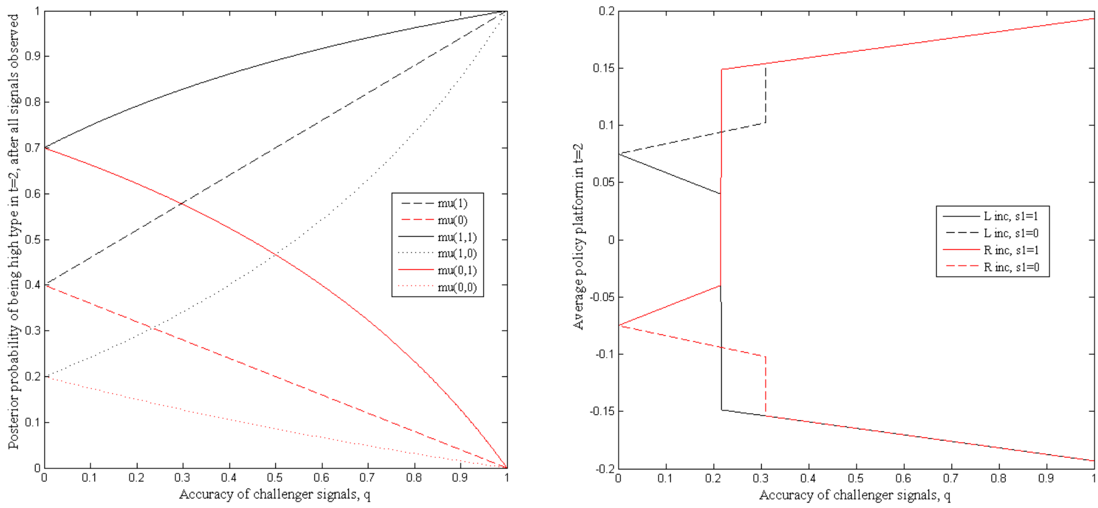

(a) Posterior beliefs in period 2 (b) Mean platform in period 2

Figure 3.1: Posteriors and mean policies

period, for each of the states (an incumbent identity and first period signal). Since I consider a case where the prior is less than $\frac{1}{2}$, the mean policy is biased against the incumbent's preferences, for all q, when the incumbent has a bad signal from the previous period. Instead, if q is large enough, the average policy in the state where the incumbent has a good signal will be biased in his favour.

Period 1

Let the continuation value of winning, conditional on producing a good signal in the first period, be V^{I1} and that conditional on producing a bad signal be V^{I0}. I also denote the continuation value of losing, conditional on the incumbent producing a good signal, as V^{C1}, and that conditional on him producing a bad signal be V^{C0}.

In an open seat election, the probabilities of winning the competence-based vote are more straightforward. L will win the full competence-based vote if he is the only player to get a good signal, and he will also get half of it if they both perform equally well.

$$p_L = \frac{1}{2} + 2\phi\alpha^v\bar{x}_1 + 2\phi\alpha^v\frac{(1-\gamma)}{\gamma}\left(\mu(1-\mu) - (1-\mu)\mu\right) = \frac{1}{2} + 2\phi\alpha^v\bar{x}_1 \qquad (3.17)$$

The expected electoral bias today from competition for competence-based votes is zero, because candidates are ex-ante symmetric in the open seat election.

The optimisation problem of the candidates needs to take into account that the continuation value of winning depends on the particular signals in this period. L

solves,

$$
\begin{aligned}
&\max_{x_L} -\alpha^v(x_R-\hat{x}^L)+p_L(x_L,x_R)\alpha^v(x_R-x_L)\\
&\quad+\mu(1-\mu)\beta\left[V^{C0}+\left(\frac{1}{2}+2\phi\alpha^v\bar{x}_1+2\phi\alpha^v\frac{(1-\gamma)}{\gamma}\right)(V^{I1}-V^{C0})\right]\\
&\quad+(1-\mu)\mu\beta\left[V^{C1}+\left(\frac{1}{2}+2\phi\alpha^v\bar{x}_1-2\phi\alpha^v\frac{(1-\gamma)}{\gamma}\right)(V^{I0}-V^{C1})\right]\\
&\quad+\mu^2\beta\left[V^{C1}+\left(\frac{1}{2}+2\phi\alpha^v\bar{x}_1\right)(V^{I1}-V^{C1})\right]\\
&\quad+(1-\mu)^2\beta\left[V^{C0}+\left(\frac{1}{2}+2\phi\alpha^v\bar{x}_1\right)(V^{I0}-V^{C0})\right]
\end{aligned}
\tag{3.18}
$$

L's first order condition looks like,

$$
\alpha^C p_L=\phi\alpha^v\left[\alpha^C(x_R-x_L)+\beta\left(\mu(V^{I1}-V^{C1})+(1-\mu)(V^{I0}-V^{C0})\right)\right] \tag{3.19}
$$

One can proceed in a similar manner to obtain R's first order condition, and combining the two expressions yields the solution in terms of the mean platform and degree of polarisation.

$$
\bar{x}_1=0 \text{ and } \delta x_1=\frac{1}{2\phi\alpha^v}-\beta\left(\mu(V^{I1}-V^{C1})+(1-\mu)(V^{I0}-V^{C0})\right) \tag{3.20}
$$

The mean policy is zero because, as an open seat election, the electoral race is unbiased today. Since the candidates are symmetric, they have the same expected boost from competing for the competence-based voters. Polarisation will arise because of uncertainty and may be accentuated or mitigated by the dynamic effect. If the expected marginal continuation value of winning is positive then the dynamic effect will moderate polarisation to some extent, by making parties compete harder for votes. So what matters when parties choose their platform is the expected marginal continuation value of winning, where the expectation is over the two possible signals that L might produce today.

The expected marginal continuation value of winning is

$$
\begin{aligned}
&\left(\mu(V^{I1}-V^{C1})+(1-\mu)(V^{I0}-V^{C0})\right)=\mu\left(-\alpha^C(\bar{x}_2^L(1)-\hat{x}^L)+\alpha^C(\bar{x}_2^R(1)-\hat{x}^L)\right)\\
&+(1-\mu)\left(-\alpha^C(\bar{x}_2^L(0)-\hat{x}^L)+\alpha^C(\bar{x}_2^R(0)-\hat{x}^L)\right)\\
&=4\alpha^C\frac{(1-\gamma)}{\gamma}\left(\mu Pr[\tilde{\mu}^L\geq\tilde{\mu}^R|s_1^L=1]+(1-\mu)Pr[\tilde{\mu}^L\geq\tilde{\mu}^R|s_1^L=0]\right)
\end{aligned}
\tag{3.21}
$$

Now it is clearer why the distinction between the electoral bias and the incumbency advantage Δ matters. The incumbency advantage is the expected electoral advantage (from winning competence-based votes), conditional on winning the first election. The degree of polarisation will instead depend on the *unconditional* expectation of incumbency bias (where the expectation is over the prior distribution of states). When candidates choose to shift their platform in order to attract more votes, they increase both the probability of winning when they produce a good signal today, and when they produce a bad one, and they increase these probabilities by the same amount. So what matters for ideological competition is the probability of producing each of those signals, and not the probability of the signals conditional on winning. This implies that, as long as the quality distribution is skewed towards low quality, the marginal continuation value of winning will be negative.

To re-iterate, the incumbency advantage is the expected probability of a candidate winning the competence-based vote in period 2, conditional on winning the election in period 1[10].

$$\begin{aligned}\Delta =& Pr(s_1^L = 1|\text{L wins at } t = 1)Pr[\tilde{\mu}^L \geq \tilde{\mu}^R|s_1^L = 1] \\ &+ Pr(s_1^L = 0|\text{L wins at } t = 1)Pr[\tilde{\mu}^L \geq \tilde{\mu}^R|s_1^L = 0] - \frac{1}{2}\end{aligned} \tag{3.22}$$

where

$$Pr(s_1^L = 1|\text{L wins at } t = 1) = 2\mu\left(\frac{1}{2} + (1-\mu)2\phi\alpha^v\frac{(1-\gamma)}{\gamma}\right)$$

Figure (3.2) shows the two for the specific case, $\alpha^C = \alpha^v = 1$ (and again, as earlier $\mu = 0.35$, $p = 0.5$, and $\gamma = 0.8$). We can see that while for high enough q there will indeed be incumbency advantage, the marginal value of winning is always negative. From the standpoint of candidates at the start of period 1, they are more likely to produce a bad signal and expect a bias against them if they make it through to the second period.

Increasing q has a non-monotonic effect on polarisation. In the low q region, the marginal continuation value of winning is at its lowest level and small increases in q have no effect at all. This is because, for the lowest q, the winner of the competence-based votes depends only on how the incumbent performs in office, campaign signals are ignored. However, as q increases, eventually campaign signals start to have an effect. First we reach the case that a bad signal from the challenger starts to be important, and this will increase the marginal continuation value of winning. At even

[10] The empirical incumbency advantage, is the expected probability of a candidate winning the overall election vote in period 2, conditional on winning the election in period 1. In my model, because I assumed a linear-loss function for policy payoffs, the probability of winning is always $\frac{1}{2}$ for each candidate. However, as soon as we introduced some concavity to the payoff functions, the empirical incumbency advantage will be increasing in the competence-based incumbency advantage.

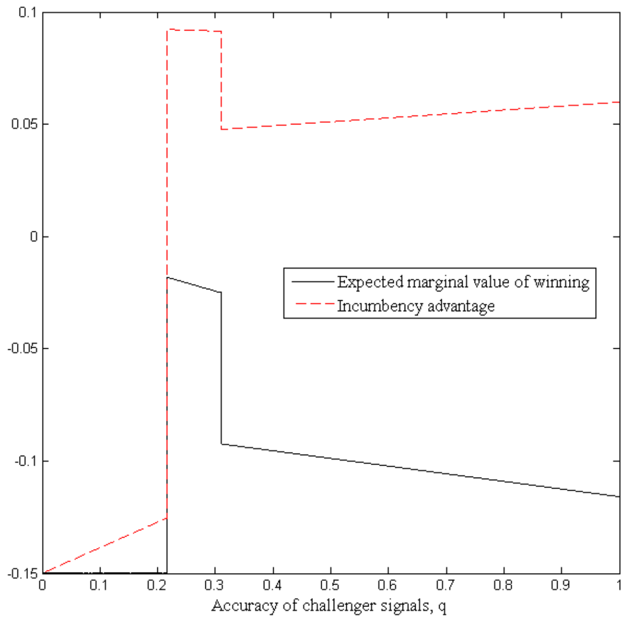

Figure 3.2: Expected marginal value of winning and incumbency advantage

higher q, challengers with good signals will beat a mixed record by the incumbent, and this reduces the marginal continuation value of winning. Also note that with higher q there will be higher correlation between the two signals of the incumbent; mixed records will be rarer, extreme ones more common. Since the prior is $\mu < \frac{1}{2}$, this means probability shifts particularly to the case where the incumbent has two bad signals. This is the reason why the marginal continuation value of winning slopes down once in the region where campaign signals matter.

In summary then, increasing q up to the region where campaign signals are taken into account by voters, can reduce the loss from winning and reduce ideological polarisation. However, further increases to q can reduce the value of winning and generate higher polarisation.

A remark on welfare implications

The example I have analysed above makes it clear that changes to the quality of information can have rather complex effects on outcomes. It is not possible to reach highly general conclusions on how welfare is affected by improvements in available information regarding political candidates. I make some observations that are informed by the example, but only qualitative in nature.

The average quality of candidates on the job will always improve when q is higher; selection is more effective. However, the effect on welfare through policy variance is not monotonic, as I already discussed.

In my example, with $\mu < \frac{1}{2}$, when q is extremely low, the welfare effect in terms of changes to policy variance are non-existent for some range of q, there may even be a gain as q moves to the intermediate q region. So it is clear that overall, starting from $q = 0$, it will be welfare-improving to increase information about challengers. Once in the intermediate q region, there is a trade-off between improving selection and disincentivising competition on policy, because the marginal value of winning drops with q. Therefore, eventually, too much information about challengers can be bad for society because it contributes to ideological extremism.

The pattern above generalises to other distributions of competence: if q is very low, improvements to information can raise welfare unambiguously both through the quality and ideological extremism channels. Eventually a trade-off develops, so that too much information may harm society by reducing the value of a win, and creating large ideological polarisation.

3.5 Conclusions

Incumbency effects are a common feature of electoral systems. My primary focus has been to analyse the relationship between incumbency-related electoral biases and ideological polarisation. My main result indicates that while any kind of incumbency-related bias pulls policies away from the median's ideal, biases that systematically favour the challenger rather than the incumbent lead to particularly extreme positions.

The static effect of any popularity advantage in elections is to shift the centre ground in favour of the more popular candidate's preferred policy. The dynamic effect contributes to divergence of policy between the parties and depends crucially on whether there is a systematic bias towards or against incumbents. Such a bias would alter the continuation value of winning and determine how strongly parties compete for the win. Biases that favour the incumbent increase the value of winning, making parties compete more fiercely, closing the policy gap. Biases that hurt the incumbent's future electoral prospects make parties less willing to compromise on policy by introducing a cost to winning.

While there may be various arguments in favour of political turnover and in some sense it may signal a well-functioning democracy, this analysis suggests a cost to alternation. Specifically, it is not that observed switching is per se bad, but if voters have a systematic preference for change this has a downside in terms of reducing the parties' willingness to compete over ideological issues. Instead of seeking out votes, they just pick policies they like and wait for the tide to turn and put them back in office. While both types of biases increase policy variance, the analysis suggests biases that favour the incumbent may well be the lesser of two evils.

I have proposed one mechanism that can generate incumbency effects. My framework is based on rational voting when there is incomplete information about the underlying competence of candidates. I find that the structure of available information matters and that incumbency disadvantage may be caused by poor selection from a pool of mostly mediocre candidates. When voters know very little about new candidates and are forced to engage in a process of trial and error for evaluating politicians' quality, they may be more likely than not to throw out incumbents. This could provide an explanation for an emerging empirical regularity concerning incumbency effects: while most advanced democracies have strong incumbency advantage, it appears that incumbency disadvantage is rather the more common feature of political office in young democracies. My results indicate that this may be due to the more developed information technologies of modern democracies, for example, a well-established media or more involved election campaigns.

Finally, note that implicit in the analysis was the possibility that differences in the incumbency effect could be explained by different quality distributions. It may be that incumbency disadvantage in some countries is due to them having worse politicians. However, I provide an alternative explanation for cross-country differences in the incumbency advantage, one based on differences in the information technologies.

Appendices

Appendix A

Chapter 1 - Proofs and Extensions

A.1 Proof of the properties of p_2 in the endogenous model of elections

From Equation (1.7), I can write the effect of a change in inequality on the probability of P winning the election in period 2,

$$
\begin{aligned}
\frac{\partial p_2}{\partial \iota_2} = \psi \Big\{ & \frac{\partial \widehat{G}^P(K_2)}{\partial \iota_2} \left(-\frac{y^P(K_2)}{y^m(K_2)} + v'(\widehat{G}^P(K_2)) \right) \\
& - \frac{\partial \widehat{G}^R(K_2)}{\partial \iota_2} \left(-\frac{y^P(K_2)}{y^m(K_2)} + v'(\widehat{G}^R(K_2)) \right) \\
& - (\widehat{G}^P(K_2) - \widehat{G}^R(K_2)) \frac{\partial (y^P(K_2)/y^m(K_2))}{\partial \iota_2} \Big\} > 0
\end{aligned}
\tag{A.1}
$$

Suppose first that K_2 is such that $y^R(K_2) > y^P(K_2)$. Using the first order conditions for $\widehat{G}^P$ and $\widehat{G}^R$ we know that the first bracketed term is zero (by P's first order condition). The second bracketed term is positive because when R is richer, R underspends as far P is concerned, so the net marginal benefit of public consumption for P is positive. The third bracket is positive because, again, P would like to set higher public consumption than R.

In addition we know that $\frac{\partial \widehat{G}_2^R}{\partial \iota} < 0$ because if R is richer, then he faces a higher price of public funds as inequality increases and spends less. Finally, when $y^R(K_2) > y^P(K_2)$ we have $\frac{\partial (y_2^P/y_2^m)}{\partial \iota} < 0$. Putting all these observations together, one concludes that higher inequality unambiguously improves the electoral chances of the P party by increasing the utility gap between its own platform (which is also that of the median voter) and that offered by R.

When $y^R(K_2) < y^P(K_2)$, that is, K_2 levels such that P overtake R and become the richer group, all the signs invert but the overall result remains the same. R will now choose higher $\widehat{G}^R$ when inequality is higher; the net marginal benefit of public consumption to P, evaluated at R's policy will be negative, because now P perceives R to be overspending; and the ratio of P's income to the mean will be increasing with the level of inequality. The result is again that higher inequality leads to a higher probability of winning.

Next consider the effect on electoral prospects of starting period 2 with a larger aggregate income, but the same level of inequality.

$$
\begin{aligned}
\frac{\partial p_2}{\partial y_2^m} = \psi \Big\{ & \frac{\partial \widehat{G}^P(K_2)}{\partial y^m(K_2)} \left(v'(\widehat{G}^P(K_2)) - \frac{y^P(K_2)}{y^m(K_2)} \right) + \frac{\partial \widehat{G}^R(K_2)}{\partial y^m(K_2)} \left(v'(\widehat{G}^R(K_2)) - \frac{y^P(K_2)}{y^m(K_2)} \right) \\
& - (\widehat{G}^P(K_2) - \widehat{G}^R(K_2)) \frac{\partial}{\partial y^m(K_2)} \left(\frac{y^P(K_2)}{y^m(K_2)} \right) \Big\}
\end{aligned}
\tag{A.2}
$$

But since the demand for public consumption has no income effect, the first two terms are zero. The final term is also zero because $(\frac{y^P(K_2)}{y^m(K_2)})$ reduces to being purely a function of $\iota(K_2)$ which we are holding fixed.

A.2 Proof of Proposition 1

I show the proof for party P, the proof for party R is analogous.

The first order condition for $\widehat{K}_2^P$ sets the marginal cost of investment equal to the marginal benefit, and the solution exists if the second order condition holds, i.e. if the MB curve is downward sloping in K_2 (the MC curve is flat).

Therefore it suffices to show that an increase in $\widehat{b}$ shifts the MB down when investment is income equalising and shifts it up when investment leads to higher inequality.

$$\begin{aligned}\frac{\partial MB}{\partial \widehat{b}} =& \psi \frac{\partial \iota_2}{\partial K_2} \frac{\partial \widehat{G}^R(K_2)}{\partial \iota_2} \left(\frac{y^P(K_2)}{y^m(K_2)} - v'(\widehat{G}^R(K_2)) \right) \\ &+ (\widehat{\tau}^P(K_2) - \widehat{\tau}^R(K_2)) y^P(K_2) \left(\frac{y_K^m(K_2)}{y^m(K_2)} - \frac{y_K^P(K_2)}{y^P(K_2)} \right)\end{aligned} \tag{A.3}$$

Suppose P is the poorer (richer) group at the equilibrium level K_2. Then, R's public consumption choice is decreasing (increasing) in ι_2 and the first bracket is negative (positive) because P-types perceive R to be underspending (overspending). In any case the product of those two terms is positive, so overall the first term in the sum is positive when investment raises inequality, and negative when it is income-equalising. For the difference in tax rates of the two groups we have $\widehat{\tau}^P(K_2) - \widehat{\tau}^R(K_2) > (<) 0$, since P wants more (less) public consumption tomorrow than R, and the last bracket is positive when aggregate income grows at a faster rate with respect to public capital than the income of the P group. Therefore the second term is also positive when public capital increases inequality and negative with it is income equalising.

In summary,

$$\frac{\partial MB}{\partial \widehat{b}} > 0 \iff \frac{\partial \iota_2}{\partial K_2} > 0 \tag{A.4}$$

A.3 Proof of Proposition 2

Suppose that, in addition to the policy motivation, when a party holds office it receives a rent r. In period 2 this has no effect because with commitment to platforms, the rent is already realised when the party is choosing its policy. It will however play a role in the first period choice of investment. In the condition for optimal investment, there is an additional bonus to the MB of increasing the party's winning probability. In the case of party P the FOC is:

$$\frac{1}{\beta}\frac{y^P(K_1)}{y^m(K_1)} = y^P_K(K_2)\left(p_2(1-\widehat{\tau}^P_2) + (1-p_2)(1-\widehat{\tau}^R_2)\right)$$
$$+ y^m_K(K_2)\left(p_2\widehat{\tau}^P_2 + (1-p_2)\widehat{\tau}^R_2\right)\frac{y^P(K_2)}{y^m(K_2)}$$
$$+ \frac{\partial \iota_2}{\partial K_2}\left[\frac{\partial p_2}{\partial \iota_2}\left(\Delta^P(\widehat{G}^P_2, \widehat{G}^R_2) + \underbrace{r}_{rent}\right) + (1-p_2)\frac{\partial \widehat{G}^R_2}{\partial \iota_2}\left(-\frac{y^P(K_2)}{y^m(K_2)} + v'(\widehat{G}^R_2)\right)\right] \tag{A.5}$$

Therefore,

$$\frac{\partial MB}{\partial r} = \frac{\partial \iota_2}{\partial K_2}\frac{\partial p_2}{\partial \iota_2} \tag{A.6}$$

The probability of party P winning always increases with inequality, because it makes the perceived difference between the two parties' platforms for the median voter larger (and he always prefers the P platform). Hence, higher rents will increase the incentive to invest when the investment increases inequality: the P party will exploit the fact that high inequality gives it an (endogenous) electoral advantage tomorrow. For R the effect will be the opposite.

The comparative static with respect to electoral uncertainty is also straightforward to find.

We find

$$\frac{\partial MB}{\partial \psi} = \frac{\partial p_2}{\partial \psi}(\widehat{\tau}^P_2 - \widehat{\tau}^R_2)y^P(K_2)\left[\frac{y^m_K(K_2)}{y^m(K_2)} - \frac{y^P_K(K_2)}{y^P(K_2)}\right]$$
$$\frac{\partial \iota_2}{\partial K_2}\left[\frac{\partial^2 p_2}{\partial \iota_2 \partial \psi}\Delta^P(\widehat{G}^P_2, \widehat{G}^R_2) - \frac{\partial p_2}{\partial \psi}\left(v'(\widehat{G}^R_2) - \frac{y^P(K_2)}{y^m(K_2)}\right)\frac{\partial \widehat{G}^R_2}{\partial \iota_2}\right] \tag{A.7}$$

The sign for this depends only on whether the investment is income-equalising or not. Consider incentives for the party P. The (equilibrium) probability of winning is increasing with ψ. The tax rate difference is positive if it is the poorer group, in which case the first term in the sum is positive overall if the investment increases inequality. Instead if it is the poorer the tax rate difference is negative and the term is positive overall if P's income grows faster than the aggregate income, which again means higher inequality. The cross-partial derivative of p_2 is positive, the more sensitive voters are to policies the more P's chances improve with the level of inequality because voters will give more weight to the growing difference in utility between the two platforms. The net marginal utility of public consumption at R's policy is positive (negative) when P is the poorer (richer) group, and $\frac{\partial \widehat{G}^R_2}{\partial \iota_2}$ is negative (positive). Therefore overall the expression in the second line of the equation is positive if and only if investment increases inequality, which goes in the same direction as the overall effect on the first line.

A.4 Generalisation of the distribution of popularity shocks

If P is the incumbent in period 1, it anticipates a probability p_2 of staying in power of

$$\begin{aligned} p_2 = Pr\Big\{&\phi(1-\widehat{\tau}^P(K_2))y^P(K_2) + v(\widehat{G}^P(K_2)) + \epsilon_2 \\ &\geq \phi(1-\widehat{\tau}^R(K_2))y^P(K_2) + v(\widehat{G}^R(K_2))\Big\} \end{aligned} \tag{A.8}$$

For a general distribution of shocks given by the CDF, $F(x)$, with domain over $(-\infty, +\infty)$, which we assume to have a bias so that $E(\epsilon_t) = b$, we have

$$\begin{aligned} p_2 =& 1 - F\left(\phi(1-\widehat{\tau}^R(K_2))y^P(K_2) + v(\widehat{G}^R(K_2)) - \phi(1-\widehat{\tau}^P(K_2))y^P(K_2) - v(\widehat{G}^P(K_2))\right) \\ =& 1 - F(-\Delta^P(\widehat{G}^P(K_2), \widehat{G}^R(K_2))) \end{aligned} \tag{A.9}$$

When deciding how much to invest in period 1, party P will consider how public capital, through its effect on inequality, will influence future electoral prospects. This term appeared in equation (1.11). However, as we can see from equation (1.8), when the shock is uniformly distributed P's probability of re-election is separable in the level of inequality and the incumbency bias, which we have just explained above is due to the linearity of the CDF. For distributions for which $F''(x) \neq 0$, changing the size of the incumbency effect will affect the MB from investment through the electoral effect.

To illustrate how results may change, I focus on another familiar distribution: let $\epsilon_t \sim N(b, \sigma^2)$. Again let the bias be linked to incumbency, according to $b = \widehat{b}I\{\tilde{J}_{t-1} = P\} - \widehat{b}I\{\tilde{J}_{t-1} = R\}$.

The normal distribution has a CDF that is concave below $\frac{1}{2}$ and convex above it. This means that when the election is a close race, in equilibrium, the returns to changing the party platform are high, but if one of the parties has a large net advantage in the election, the returns decline. Now there is an interaction between the degree of incumbency advantage and the returns to investment in terms of improved electoral advantage.

The probability of P winning tomorrow, if it is today's incumbent is,

$$\begin{aligned} p_2 =& Pr\left[\epsilon_2 \geq -\Delta^P(\widehat{G}_2^P, \widehat{G}_2^R)\right] = 1 - \Phi\left(\frac{-\Delta^P(\widehat{G}_2^P, \widehat{G}_2^R) - \widehat{b}}{\sigma}\right) \\ =& \Phi\left(\frac{\Delta^P(\widehat{G}_2^P, \widehat{G}_2^R) + \widehat{b}}{\sigma}\right) \end{aligned} \tag{A.10}$$

It is safe to assume that when the P party is in power, it will always have a net advantage in the race since both the incumbency bias and the fact that the median voter always prefers its platform to R's, contribute to making the P party more likely to win. The magnitude of incumbency advantage now affects the marginal benefit of investment through the electoral effect according to

$$\frac{\partial p_2}{\partial \widehat{b} \partial \iota_2} = \frac{1}{\sigma^2} \frac{\partial \Delta^P(\widehat{G}_2^P, \widehat{G}_2^R)}{\partial \iota} \Phi'' \left(\frac{\Delta^P(\widehat{G}_2^P, \widehat{G}_2^R) + \widehat{b}}{\sigma} \right) \tag{A.11}$$

Since in equilibrium $p_2 > \frac{1}{2}$ when P is the incumbent, then the CDF is concave in this region and $\Phi''(x) < 0$. This means that for higher $\widehat{b}$ the effect of changing inequality on the probability of winning is weaker and P will have less incentive to choose high levels of inequality. This effect therefore goes in the opposite direction to the other effects influencing investment choices, and if it is strong enough, it could reverse the results from Proposition 1.

The intuition is as follows. P can maintain a high probability of survival in office by maintaining a high level of inequality, because with higher inequality the median voter experiences a larger utility loss from electing R rather than P. If the effectiveness of such policies at securing votes dies off as a party takes a bigger lead, then when party P has a larger incumbency advantage it will see less of a need to maintain high inequality as a way of securing its political rule.

For R, the direction of the effect of incumbency advantage on the electoral term of the marginal benefit of investment, depends on whether in equilibrium R manages to lead the race when it is the incumbent. For relatively small $\widehat{b}$, R will face a probability of survival smaller than $\frac{1}{2}$ and increases in $\widehat{b}$ will raise the returns to changing the party platform, or, in R's case the return to reducing inequality before the period 2 election. So again it is possible that the electoral effect runs counter to the strategic and price effects.

Once $\widehat{b}$ is so large that in equilibrium R has a net electoral advantage whenever it is the incumbent, the results above reverse. In this case, further increases in $\widehat{b}$ reduce the slope of the CDF, i.e. the returns to reducing inequality. Now, R will be less concerned about maintaining low inequality.

I have analysed how the results change when the popularity shocks follow a normal distribution with mean b. I find that for the P party the electoral effect always opposes the result from Proposition 1. For party R the electoral effect opposes the result from Proposition 1 whenever the equilibrium probability of winning for party P is less than $\frac{1}{2}$, but once this probability is higher than $\frac{1}{2}$, it reinforces the the result from Proposition 1.

More generally it is now clear that the results depend on whether the CDF is concave or convex in the relevant region of electoral probabilities. This will determine whether shifts in the CDF as the incumbency advantage rises, raise or lower the returns to changes in inequality in terms of boosting the party's electoral prospects.

A.5 The central planner's infinite horizon problem

The programme to be solved by the central planner in this case is

$$\begin{aligned}\max_{\{\tau_t, G_t, K_t'\}_{t=0}^{\infty}} &\sum_{t=0}^{\infty} \beta^t \left[\phi(1-\tau_t)y^m(K_t) + v(G_t)\right] \\ &\text{s.t. } \tau_t y^m(K_t) = G_t + K_t' - (1-d)K_t \\ &K_{t+1} = K_t',\ K_0 \text{ given.}\end{aligned} \tag{A.12}$$

The problem is recursive, with a single state variable K_t, so that the value function must satisfy the following functional equation.

$$V(K_t) = \phi(1-\tau_t)y^m(K_t) + v(G_t) + \beta V(K_t') \tag{A.13}$$

subject to the constraint $\tau_t y^m(K_t) = G_t + K_t' - (1-d)K_t$.

The optimal public consumption policy is constant and given by the condition $\phi = v'(\bar{G}^{cp})$. The first order condition for investment is

$$\phi = \beta V'(K_t') \tag{A.14}$$

But this indicates that the optimal choice of next period capital doesn't depend on the current level of capital. The system will jump immediately into a steady state with capital $\bar{K}^{cp}$. I can further pin down this steady state by considering the problem of the central planner choosing today, taking as given his future self's choices.

$$\begin{aligned}\max_{\tau_t, K_t'} &\phi(1-\tau_t)y^m(K_t) + v(\bar{G}^{cp}) \\ &+ \beta\left[\left(1 - \frac{\bar{G}^{cp} + \bar{K} - (1-d)K_t'}{y^m(K_t')}\right) y^m(K_t') + v(\bar{G}^{cp}) + \beta V(\bar{K}^{cp})\right]\end{aligned} \tag{A.15}$$

subject to the constraint $\tau_t y^m(K_t) = \bar{G}^{cp} + K_t' - (1-d)K_t$. Which has the following first order condition for investment

$$\phi = \beta\left(1 - d + y_K^m(K_t')\right) \tag{A.16}$$

The only difference with respect to the two-period game is that, because investment will continue in the future, the marginal benefit of investment today includes the value of undepreciated capital tomorrow.

Therefore, in the dynamic model there will be a constant policy plan $\{\bar{G}^{cp}, \bar{K}^{cp}\}$, where the investment choice satisfies the condition

$$\phi = \beta\left(1 - d + y^m(\bar{K}^{cp})\right) \tag{A.17}$$

A.6 Pure redistribution - welfare analysis

I consider an environment that I revisit in the dynamic analysis of Section 1.5. Suppose that aggregate income is fixed with respect to public capital, $y^m(K) = y$ for all $K \geq 0$. Public capital is income-equalising, specifically $y^P(K) = a + gK$ where $a, g > 0$.

The value of public consumption is $v(G) = \gamma G - \frac{G^2}{2}$, where $\gamma > 0$. This implies the following public consumption choices for the two groups,

$$G^J(K) = \gamma - \phi \frac{y^J(K)}{y}$$

As an example let me take the following parameter configuration: $a = 50$, $g = 50$, $y = 200$, $\gamma = 50$, $\phi = 50$, $\beta = 0.99$, $d = 1$, $\alpha = 0.5$. Also for simplicity let the probability of re-election be exogenous, and set at $p = 0.35$. Finally, suppose initially there is no public capital, $K_1 = 0$.

I solve for consumption and investment choices in both the strategic environment and in the central planner problem. The central planner chooses intermediate levels of public consumption in every instance; for example, in the first period she chooses $\bar{G}^{cp} = 150$ while P chooses $\bar{G}^P(0) = 187.5$ and R chooses $\bar{G}^P(0) = 112.5$. Both groups choose positive levels of public capital for the second period: $\widehat{K}_2^P = 5.45$ and $\widehat{K}_2^R = 0.15$. In contrast, since the central planner only values public capital for its effect on aggregate income, she does not invest at all, $\widehat{K}_2^{cp} = 0$.

Why is even the R group willing to invest more than a central planner, even though public capital hurts its private income and does no good in terms of raising aggregate income? This is driven by the strategic role of investment: R invests because it is concerned about the possibility of P taking over in office tomorrow and over-spending on public consumption because it faces a very low price of public funds. R brings down inequality to reduce this conflict.

The central planner has no such concern, any redistribution concerns can be addressed tomorrow, she does not need to pre-commit redistribution in advance through investment. Since this investment is costly, it is clear that the strategic environment results in inefficiency, resources today are used up without any increase in available resources tomorrow.

Are the outcomes under delegation a Pareto improvement on the strategic outcomes? I compute the welfare of each of the groups when they start in power, but face the threat of turnover[1], and compare it to their welfare under the central planner policies. For this parameter configuration, I find that the R group would prefer to delegate to a central planner, while the P group would prefer not to, so delegation is not a Pareto improvement. These preferences hold whether I compare outcomes

[1]The value of the game for J when it starts off as incumbent is

$$\begin{aligned}&\phi(1-\widehat{\tau}^J)y^J(0) + v(\widehat{G}^J(0))\\&+\beta\left[p\left(\phi(1-\widehat{\tau}^J)y^J(\widehat{K}^J) + v(\widehat{G}^J(\widehat{K}^J))\right) + (1-p)\left(\phi(1-\widehat{\tau}^{-J})y^J(\widehat{K}^J) + v(\widehat{G}^{-J}(\widehat{K}^J))\right)\right]\end{aligned}$$

under delegation to expected welfare in the strategic environment conditional on being the initial incumbent and if the initial election was an open seat election[2].

How is welfare affected by incumbency advantage? If a group finds itself in power, its welfare increases with the degree of incumbency advantage: holding onto power has value. However, if I allow for the identity of the first period incumbent to be determined in a fair open seat election, then both groups have higher ex-ante welfare when the probability of turnover is higher. Again a higher threat of losing office makes rivals more willing to avoid severe conflict in future, and both groups gain from the lower levels of inequality.

A.7 Proof of Lemma 3

I adjust the proof from Fudenberg and Tirole 1991, to a situation where the order of play is stochastic.

The set of players is now $\{0, 1P, 1R, 2P, 2R....\}$, which is countable. I have a game of incomplete information because nature picks, after every move I_t, whether it is player $(t+1)P$ who plays or player $(t+1)R$. Therefore the continuation value of an action is an expectation that can be computed if I fix players' strategies, so that player tJ plays $s^{tJ}(K_t)$. Specifically let the strategy be the capital level for the next period.

Fix Markov strategies for players $(t+1)J$, $(t+2)J$, etc and let

$$\bar{w}_t^J(K_{t+1}) \equiv E_t w_t^J(K_{t+1}, s^{(t+1)\tilde{J}}(K_{t+1}), s^{(t+2)\tilde{J}}(f(s^{(t+1)\tilde{J}}(K_{t+1}))), ...)$$

where the expectation is over the type of the decision maker $\tilde{J}$, which follows the stochastic process described in section 2 (that is, the decision maker from period t is again decision maker in period $t+1$ with probability p and the identity of the decision maker switches with probability $1-p$).

The rest of the proof is standard. Consider two states K and $\tilde{K}$ and denote the optimal action at each state by K' and $\tilde{K}'$ respectively. By optimality we have,

$$u^J(K, K') + \bar{w}_t^J(K') \geq u^J(K, \tilde{K}') + \bar{w}_t^J(\tilde{K}') \tag{A.18}$$

and similarly,

$$u^J(\tilde{K}', \tilde{K}') + \bar{w}_t^J(\tilde{K}') \geq u^J(\tilde{K}, K') + \bar{w}_t^J(K') \tag{A.19}$$

[2]The value of the game to group J when the period 1 incumbent is decided in a fair open seat election is

$$\frac{1}{2}\Big\{\phi(1-\hat{\tau}^J)y^J(0) + v(\widehat{G}^J(0)) + \beta\Big\{p\left(\phi(1-\hat{\tau}^J)y^J(\widehat{K}^J) + v(\widehat{G}^J(\widehat{K}^J))\right)$$
$$+ (1-p)\left(\phi(1-\hat{\tau}^{-J})y^J(\widehat{K}^J) + v(\widehat{G}^{-J}(\widehat{K}^J))\right)\Big\}\Big\} + \frac{1}{2}\Big\{\phi(1-\hat{\tau}^{-J})y^J(0) + v(\widehat{G}^{-J}(0))$$
$$+ \beta\Big\{p\left(\phi(1-\hat{\tau}^{-J})y^J(\widehat{K}^{-J}) + v(\widehat{G}^{-J}(\widehat{K}^{-J}))\right) + (1-p)\left(\phi(1-\hat{\tau}^J)y^J(\widehat{K}^{-J}) + v(\widehat{G}^J(\widehat{K}^{-J}))\right)\Big\}\Big\}$$

Adding the two expressions eliminates the continuation value terms and leaves,

$$u^J(K,K') + u^J(\tilde{K}',\tilde{K}') - u^J(K,\tilde{K}') - u^J(\tilde{K},K') \geq 0 \tag{A.20}$$

which can be rewritten as,

$$\int_{K'}^{\tilde{K}'}\int_{K}^{\tilde{K}} \frac{\partial u}{\partial x \partial y} dx dy \geq 0 \tag{A.21}$$

Hence, if $\tilde{K} > K$, then it follows that $\tilde{K}' > K'$ if (CS^+) holds and $\tilde{K}' < K'$ if (CS^-) holds.

A.8 Proof of Lemma 4

I first consider the case in which P is the only group in the world and its income is given by $y^P(K) = a + gK$ (note that now aggregate income must equal y^P because there is no other group).

P will solve

$$\max_{\{G_t,K_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \left[\phi\left(1 - \frac{G_t + K_{t+1} - (1-d)K_t}{a + gK_t}\right)(a + gK_t) + v(G_t)\right] \tag{A.22}$$

Since the government now faces a constant price of public expenditures it will choose a constant level of public consumption give by

$$\phi = v'(\bar{G}^P)$$

Additionally, since the marginal benefit curve for next period public capital doesn't depend on the current level, the government will either jump immediately to a (finite) steady state, never invest, or we might find that there is no solution.

Suppose that an interior solution exists, such that the government jumps to the steady state $\bar{K}$. Then in equilibrium, the value of the game in the steady state is,

$$V(\bar{K}) = \sum_{t=0}^{\infty} \beta^t \left[\phi\left(1 - \frac{\bar{G}^P + d\bar{K}}{a + g\bar{K}}\right)(a + g\bar{K}) + v(\bar{G}^P)\right] \tag{A.23}$$

If this is the equilibrium, then there should be no incentive to deviate from level $\bar{K}$ in any one period. A one-shot deviation to capital level K' will give P a value of

$$\begin{aligned} V^{dev} =& \phi\left(1 - \frac{\bar{G}^P + K' - (1-d)\bar{K}}{a + g\bar{K}}\right)(a + g\bar{K}) + v(\bar{G}^P) \\ &+ \beta\left[\phi\left(1 - \frac{\bar{G}^P + \bar{K} - (1-d)K'}{a + b\bar{K}'}\right)(a + g\bar{K}') + v(\bar{G}^P)\right] + \beta^2 V(\bar{K}) \end{aligned} \tag{A.24}$$

A deviation is profitable if $V^{dev} > V(\bar{K})$, that is if,

$$\phi(\bar{K} - K') + \beta\left[\phi\left(g(K' - \bar{K}) - (1-d)(\bar{K} - K')\right)\right] \tag{A.25}$$

Whether this condition holds doesn't depend on the particular deviation considered, other than whether they are to higher or lower levels than $\bar{K}$. Deviations to higher levels will be profitable whenever $\beta g > (1 - \beta(1-d))$, and so in this case, there is no equilibrium because P would always prefer to deviate to even higher K. If $\beta g < (1-\beta(1-d))$, then P would like to deviate to lower K and so the equilibrium involves never investing. Only in the knife-edge case of $\beta g = (1 - \beta(1-d))$ can there be an interior solution.

Turn to the case where only R exists, and has income $y^R(K) = a^R + g^R K$ where, as earlier $g^R < 0$, R becomes poorer with higher levels of K. Again since the price of public expenditures is constant, G^R will be constant, and public capital must either jump to a steady state, remain at zero, or there is no solution.

I can follow exactly the same approach as above and get the following condition for R to want to deviate to higher levels of public capital: $\beta g^R > (1 - \beta(1-d))$. Since $g^R < 0$ this cannot hold, and neither can the condition for indifference which allows for interior solutions. R will never invest if there is no P group.

Might R want to invest if the poor group exists but never takes power? I use a similar approach, looking at one shot deviations from an assumed equilibrium profile of strategies. In this case, let me suppose that the world starts off with no public capital and R would choose $K_{t+1} = 0$ in every period from tomorrow onwards. Could he benefit from deviating to a positive public capital choice K_1? This would generate the payoff

$$\begin{aligned} V^{dev} =& \phi\left(1 - \frac{G^R(0) + K'}{y}\right) a^R + v(G^R(0)) \\ &+ \beta\left[\phi\left(1 - \frac{G^R(K') - (1-d)K'}{y}\right)(a^R + g^R K') + v(G^R(K'))\right] \\ &+ \frac{\beta^2}{1-\beta}\left[\phi\left(1 - \frac{G^R(0)}{y}\right) a^R + v(G^R(0))\right] \end{aligned} \tag{A.26}$$

One can see whether the value of the game is increasing with investment, by checking the first derivative of V^{dev} with respect to K', evaluated at $K' \geq 0$:

$$\begin{aligned} \left.\frac{\partial V^{dev}}{\partial K'}\right|_{K' \geq 0} =& -\phi\frac{a^R}{y} + \beta\Bigg\{\phi(1-d)\frac{\overbrace{a^R + g^R K'}^{<a^R}}{y} + \beta\phi\left(1 - \frac{G^R(K') - (1-d)K'}{y}\right)\overbrace{g^R}^{<0} \\ &+ \frac{\partial G^R(K')}{\partial K}\underbrace{\left(v'(G^R(K')) - \phi\frac{a^R + g^R K'}{y}\right)}_{=0 \text{ by FOC of G}}\Bigg\} < 0 \end{aligned} \tag{A.27}$$

So R would always lose from raising it's choice of next period capital (and gain from lowering it), and therefore, starting from zero initial capital, with R as a dictator, there will never be any investment.

A.9 Linear-quadratic model: solving for the value functions

The Bellman equation for P in the state where he is incumbent, is:

$$
\begin{aligned}
A_{VI}+&B_{VI}K+\frac{C_{VI}}{2}K^2 = \\
&\phi\left(1-\frac{\left(\gamma-\phi\frac{(a+gK)}{y}\right)+\left(\frac{\phi(a+gK)}{\beta y C_V}-\frac{B_V}{C_V}\right)-(1-d)K}{y}\right)(a+gK) \\
&+\gamma\left(\gamma-\phi\frac{(a+gK)}{y}\right)-\frac{1}{2}\left(\gamma-\phi\frac{(a+gK)}{y}\right)^2 \\
&\beta\left[A_V+B_V\left(\frac{\phi(a+gK)}{\beta y C_V}-\frac{B_V}{C_V}\right)+\frac{C_V}{2}\left(\frac{\phi(a+gK)}{\beta y C_V}-\frac{B_V}{C_V}\right)^2\right]
\end{aligned} \tag{A.28}
$$

The Bellman equation for P in the state where he is the challenger is:

$$
\begin{aligned}
A_{VI}+&B_{VI}K+\frac{C_{VI}}{2}K^2 = \\
&\phi\left(1-\frac{\left(\gamma-\phi\frac{(a^R+g^RK)}{y}\right)+\left(\frac{\phi(a^R+g^RK)}{\beta y C_W}-\frac{B_W}{C_W}\right)-(1-d)K}{y}\right)(a+gK) \\
&+\gamma\left(\gamma-\phi\frac{(a^R+g^RK)}{y}\right)-\frac{1}{2}\left(\gamma-\phi\frac{(a^R+g^RK)}{y}\right)^2 \\
&\beta\Bigg\{(pA_{VC}+(1-p)A_{VI})+(pB_{VC}+(1-p)B_{VI})\left(\frac{\phi(a^R+g^RK)}{\beta y C_W}-\frac{B_W}{C_W}\right) \\
&+\beta\frac{(pC_{VC}+(1-p)C_{VI})}{2}\left(\frac{\phi(a^R+g^RK)}{\beta y C_W}-\frac{B_W}{C_W}\right)^2\Bigg\}
\end{aligned} \tag{A.29}
$$

We can similarly write down the Bellman equations for player R. When he is an incumbent we have:

$$
\begin{aligned}
A_{WI}+&B_{WI}K+\frac{C_{WI}}{2}K^2 = \\
&\phi\left(1-\frac{\left(\gamma-\phi\frac{(a^R+g^RK)}{y}\right)+\left(\frac{\phi(a^R+g^RK)}{\beta y C_W}-\frac{B_W}{C_W}\right)-(1-d)K}{y}\right)(a^R+g^RK) \\
&+\gamma\left(\gamma-\phi\frac{(a^R+g^RK)}{y}\right)-\frac{1}{2}\left(\gamma-\phi\frac{(a^R+g^RK)}{y}\right)^2 \\
&\beta\left[A_W+B_W\left(\frac{\phi(a^R+g^RK)}{\beta y C_W}-\frac{B_W}{C_W}\right)+\frac{C_W}{2}\left(\frac{\phi(a^R+g^RK)}{\beta y C_W}-\frac{B_W}{C_W}\right)^2\right]
\end{aligned} \tag{A.30}
$$

When he is a challenger:

$$
\begin{aligned}
A_{WI}+&B_{WI}K + \frac{C_{WI}}{2}K^2 = \\
&\phi \left(1 - \frac{\left(\gamma - \phi\frac{(a+gK)}{y}\right) + \left(\frac{\phi(a+gK)}{\beta y C_V} - \frac{B_V}{C_V}\right) - (1-d)K}{y} \right) (a^R + g^R K) \\
&+ \gamma \left(\gamma - \phi\frac{(a+gK)}{y}\right) - \frac{1}{2}\left(\gamma - \phi\frac{(a+gK)}{y}\right)^2 \\
&\beta \Bigg\{ (pA_{WC} + (1-p)A_{WI}) + (pB_{WC} + (1-p)B_{WI}) \left(\frac{\phi(a+gK)}{\beta y C_V} - \frac{B_V}{C_V}\right) \\
&+ \beta\frac{(pC_{WC} + (1-p)C_{WI})}{2}\left(\frac{\phi(a+gK)}{\beta y C_V} - \frac{B_V}{C_V}\right)^2 \Bigg\}
\end{aligned}
\tag{A.31}
$$

From the Bellman equations (A.28)-(A.31), applying the method of undetermined coefficients, I obtain twelve equations in twelve unknowns, that provide the implicit solution for the value functions in the MPE.

$$A_{VI} = \phi a - \frac{\phi\gamma a}{y} + \frac{\phi^2 a^2}{2y^2} - \frac{\phi^2 a^2}{2\beta y^2 C_V} + \frac{\phi a B_V}{y C_V} + \frac{\gamma^2}{2} + \beta A_V - \frac{\beta B_V^2}{2C_V} \tag{A.32}$$

$$B_{VI} = \frac{\phi^2 ag}{y^2} - \frac{\phi^2 ag}{\beta y^2 C_V} + \frac{\phi(1-d)a}{y} + \phi g - \frac{\phi\gamma g}{y} + \frac{\phi g B_V}{y C_V} \tag{A.33}$$

$$C_{VI} = \frac{\phi^2 g^2}{2y^2} - \frac{\phi^2 g^2}{2\beta y^2 C_V} + \frac{\phi(1-d)g}{y} \tag{A.34}$$

$$A_{WI} = \phi a^R - \frac{\phi\gamma a^R}{y} + \frac{\phi^2 (a^R)^2}{2y^2} - \frac{\phi^2 (a^R)^2}{2\beta y^2 C_W} + \frac{\phi a^R B_W}{y C_W} + \frac{\gamma^2}{2} + \beta A_W - \frac{\beta B_W^2}{2C_W} \tag{A.35}$$

$$B_{WI} = \frac{\phi^2 a^R g^R}{y^2} - \frac{\phi^2 a^R g^R}{\beta y^2 C_W} + \frac{\phi(1-d)a^R}{y} + \phi g^R - \frac{\phi\gamma g^R}{y} + \frac{\phi g^R B_W}{y C_W} \tag{A.36}$$

$$C_{WI} = \frac{\phi^2 (g^R)^2}{2y^2} - \frac{\phi^2 (g^R)^2}{2\beta y^2 C_W} + \frac{\phi(1-d)g^R}{y} \tag{A.37}$$

$$
\begin{aligned}
A_{VC} =& \phi a - \frac{\phi\gamma a}{y} + \frac{\phi^2 a a^R}{y^2} - \frac{\phi^2 a a^R}{\beta y^2 C_W} + \frac{\phi a B_W}{y C_W} + \frac{\gamma^2}{2} - \frac{\phi^2 (a^R)^2}{2y^2} \\
&+ \beta(pA_{VC} + (1-p)A_{VI}) + \beta(pB_{VC} + (1-p)B_{VI}) \left(\frac{\phi a^R}{\beta y C_W} - \frac{B_W}{C_W} \right) \\
&+ \frac{\beta}{2}(pC_{VC} + (1-p)C_{VI}) \left(\frac{\phi^2 (a^R)^2}{\beta^2 y^2 C_W^2} + \frac{B_W^2}{C_W^2} - \frac{2\phi a^R B_W}{\beta y C_W^2} \right)
\end{aligned}
\tag{A.38}
$$

$$
\begin{aligned}
B_{VC} =& \frac{\phi^2 a g^R}{y^2} - \frac{\phi^2 a g^R}{\beta y^2 C_W} + \frac{\phi(1-d)a}{y} + \phi g - \frac{\phi\gamma g}{y} + \frac{\phi a^R b}{y^2} - \frac{\phi a^R g}{\beta y^2 C_W} \\
&+ \frac{\phi g B_W}{y C_W} - \frac{\phi^2 a^R g^R}{y^2} + (pB_{VC} + (1-p)B_{VI}) \frac{\phi g^R}{y C_W} \\
&(pC_{VC} + (1-p)C_{VI}) \left(\frac{\phi^2 a^R g^R}{\beta y^2 C_W^2} - \frac{\phi g^R B_W}{y C_W^2} \right)
\end{aligned}
\tag{A.39}
$$

$$
C_{VC} = \frac{\phi^2 g g^R}{y^2} - \frac{\phi^2 g g^R}{\beta y^2 C_W} + \frac{\phi(1-d)g}{y} - \frac{\phi^2 (g^R)^2}{2y^2} + \frac{(pC_{VC} + (1-d)C_{VI})}{2} \frac{\phi^2 (g^R)^2}{\beta y^2 C_W^2}
\tag{A.40}
$$

$$
\begin{aligned}
A_{WC} =& \phi a^R - \frac{\phi\gamma a^R}{y} + \frac{\phi^2 a a^R}{y^2} - \frac{\phi^2 a a^R}{\beta y^2 C_W} + \frac{\phi a^R B_V}{y C_V} + \frac{\gamma^2}{2} - \frac{\phi^2 a^2}{2y^2} \\
&+ \beta(pA_{WC} + (1-p)A_{WI}) + \beta(pB_{WC} + (1-p)B_{WI}) \left(\frac{\phi a}{\beta y C_V} - \frac{B_V}{C_V} \right) \\
&+ \frac{\beta}{2}(pC_{WC} + (1-p)C_{WI}) \left(\frac{\phi^2 a^2}{\beta^2 y^2 C_V^2} + \frac{B_V^2}{C_V^2} - \frac{2\phi a B_V}{\beta y C_V^2} \right)
\end{aligned}
\tag{A.41}
$$

$$
\begin{aligned}
B_{WC} =& \frac{\phi_V}{+} \frac{\phi(1-d)a^R}{y} + \phi g^R - \frac{\phi\gamma g^R}{y} + \frac{\phi a g^R}{y^2} - \frac{\phi a g^R}{\beta y^2 C_V} \\
&+ \frac{\phi g^R B_V}{y C_V} - \frac{\phi^2 a g}{y^2} + (pB_{WC} + (1-p)B_{WI}) \frac{\phi b}{y C_V} \\
&(pC_{WC} + (1-p)C_{WI}) \left(\frac{\phi^2 a g}{\beta y^2 C_V^2} - \frac{\phi g B_V}{y C_V^2} \right)
\end{aligned}
\tag{A.42}
$$

$$
C_{WC} = \frac{\phi^2 g g^R}{y^2} - \frac{\phi^2 g g^R}{\beta y^2 C_V} + \frac{\phi(1-d)g^R}{y} - \frac{\phi^2 g^2}{2y^2} + \frac{(pC_{WC} + (1-d)C_{WI})}{2} \frac{\phi^2 g^2}{\beta y^2 C_V^2}
\tag{A.43}
$$

Appendix B

Chapter 2 - Proofs and Extensions

B.1 Proof of Proposition 3

My approach will be as follows. I will compute the difference between the expected payoff from banning and from allowing, denoting this as $IB\left(\tilde{y}_0, \tilde{\varepsilon}_0\right)$, the net incentive to ban. I will then show monotonicity of $IB\left(\tilde{y}_0, \tilde{\varepsilon}_0\right)$ with respect to both its arguments within three different regions of the state space. This implies a threshold solution (within each of the regions) along either dimension of the state space.

The region with $\tilde{\varepsilon}_0 \in [0, \hat{\varepsilon}]$

In the region where $\tilde{\varepsilon}_0 \leq \hat{\varepsilon}$, and given assumption 1,

$$\begin{aligned} IB\left(\tilde{\varepsilon}_0, \tilde{y}_0\right) = &\left(\rho\tilde{y}_0 - 1\right)\left(a - \tilde{\varepsilon}_0 L\right) - \left(1-\varepsilon\right)\left(\rho\tilde{y}_0 + 1 - \rho - 1\right)\left(a - \delta\tilde{\varepsilon}_0 L\right) \\ &- \varepsilon\left(\rho\tilde{y}_0 + 1 - \rho\right)\left(a - \left(\delta\tilde{\varepsilon}_0 + 1 - \delta\right)\right) L \end{aligned} \tag{B.1}$$

Monotonicity is given by the partial derivatives:

$$\frac{\partial IB\left(\tilde{\varepsilon}_0, \tilde{y}_0\right)}{\partial \tilde{\varepsilon}_0} = 1 - \rho\tilde{y}_0 - \delta\rho\left(1 - \tilde{y}_0\right) + \delta\varepsilon L > 0 \tag{B.2}$$

$$\frac{\partial IB\left(\tilde{\varepsilon}_0, \tilde{y}_0\right)}{\partial \tilde{y}_0} = \left(1-\delta\right)\rho L\left(\varepsilon - \tilde{\varepsilon}_0\right) \tag{B.3}$$

Equation B.2, along with the continuity of $IB\left(\tilde{\varepsilon}_0, \tilde{y}_0\right)$, implies that for every $\tilde{y}_0$ there is a unique threshold along the $\tilde{\varepsilon}_0$ dimension below which it is optimal to ban and above which it is optimal to allow. Call the threshold $\mu_L\left(\tilde{y}_0\right)$. The threshold will vary continuously with $\tilde{y}_0$ and therefore trace out a continuous boundary between a region of allowing (A_L) and a region of banning (B).

The thresholds in the low $\tilde{\varepsilon}_0$ region are given, for any $\tilde{y}_0$, by the following expression

$$\mu_L\left(\tilde{y}_0\right) = \hat{\varepsilon} + \frac{\left(1-\delta\right)\left[\hat{\varepsilon}\left(\varepsilon - \rho\left(1-\tilde{y}_0\right)\right) - \varepsilon\left(\rho\tilde{y}_0 + 1 - \rho\right)\right]}{1 - \rho\tilde{y}_0 + \delta\left(\varepsilon - \rho\left(1 - \tilde{y}_0\right)\right)} \tag{B.4}$$

Next I argue that if the region A_L exists it must contain the state $(0,0)$. Suppose that the threshold at the lowest policy record is below zero: $\mu_L(0) < 0$. By equation B.2 we know that this means there is no allowing in states $(0, \tilde{\varepsilon}_0)$ for $\tilde{\varepsilon}_0 \in (0, \widehat{\varepsilon})$. By equation B.3, $\frac{\partial IB(0,\tilde{\mu}_L(0))}{\partial \tilde{y}_0} > 0$. But then as we move north the incentive to ban gets stronger, and the threshold as we move north must be lower and again negative. Now we can apply the same logic as we move to higher values of $\tilde{y}_0$ and the threshold will always decrease and stay negative so that allowing will not be optimal at any state with $\tilde{\varepsilon}_0 \in [0, \widehat{\varepsilon}]$.

Will the boundary of the regions A_L and B be upward or downward sloping? I use the same arguments as above. There are three cases. (i) Suppose that $\mu_L\left(0\right) < \varepsilon$. Then by equation B.3, $\frac{\partial IB(\tilde{\mu}_L(0),0)}{\partial \tilde{y}_0} > 0$ and moving north from the point $\left(0, \mu_L\left(0\right)\right)$ the incentive to ban increases so that it must be optimal to ban north of this point. Then for $\tilde{y}_0 \gtrsim 0$ it must be that $\mu_L\left(\tilde{y}_0\right) < \mu_L\left(0\right) < \varepsilon$. Again at the state on the boundary the partial derivative in B.3 will be positive and we can iterate this argument all the way to $\tilde{y}_0 = 1$. Therefore, if $\mu_L(0) < \varepsilon$ the boundary separating the states where it is optimal to allow from those where it is optimal to ban will be downward sloping in the $(\tilde{y}_0, \tilde{\varepsilon}_0)$ state space. (ii) If $\mu_L(0) = \varepsilon$ then we have that $\frac{\partial IB(\tilde{\mu}_L(0),0)}{\partial \tilde{y}_0} = 0$ and as we move north, the indifference condition between allowing and banning will continue to hold. Then it follows that the boundary of the two regions will be exacly vertical. (iii) Finally, if $\mu_L(0) > \varepsilon$ then $\frac{\partial IB(\tilde{\mu}_L(0),0)}{\partial \tilde{y}_0} < 0$ and the incentive to ban decreases northwards. Then the threshold for $\tilde{y}_0 \gtrsim 0$ must be higher than $\mu_L(0)$ and so again higher than ε so that the incentive to ban at the boundary is again decreasing in $\tilde{y}_0$. Then the boundary must be upward sloping.

The condition on the parameters of the environment so that $\mu_L(0) < \varepsilon$ is

$$\widehat{\varepsilon}(1-\rho) - \varepsilon(2-\rho-\delta-\widehat{\varepsilon}) - \delta\varepsilon^2 < 0$$

If $\varepsilon = 0$ this condition fails and we are in case (iii) analysed above. If $2-\rho-\delta-\widehat{\varepsilon} > 0$ then the LHS of equation B.1 is monotonically decreasing in ε and therefore the condition will hold for high enough ε. If $2-\rho-\delta-\widehat{\varepsilon} > 0$ then the LHS of equation B.1 may initially increase with ε but for high enough ε the quadratic term will dominate and so again the condition will hold for high enough ε.

The region with $\tilde{\varepsilon}_0 \in [\hat{\varepsilon}, \frac{\hat{\varepsilon}}{\delta}]$

In the intermediate $\tilde{\varepsilon}_0$ region, where $\hat{\varepsilon} \leq \tilde{\varepsilon}_0 \leq \frac{\hat{\varepsilon}}{\delta}$, if the DM allows and no disaster happens the public will switch from supporting a ban to supporting allowing.

$$\begin{aligned} IB(\tilde{\varepsilon}_0, \tilde{y}_0) =& \rho\tilde{y}_0(a - \tilde{\varepsilon}_0 L) - (1-\varepsilon)(\rho\tilde{y}_0 + 1 - \rho - 1)(a - \delta\tilde{\varepsilon}_0 L) \\ &- \varepsilon(\rho\tilde{y}_0 + 1 - \rho)(a - (\delta\tilde{\varepsilon}_0 + 1 - \delta)L) \end{aligned} \tag{B.5}$$

Monotonicity is given by the partial derivatives,

$$\frac{\partial IB(\tilde{\varepsilon}_0, \tilde{y}_0)}{\partial \tilde{\varepsilon}_0} = L[\delta(\varepsilon - \rho) - \rho\tilde{y}_0(1-\delta)] \tag{B.6}$$

$$\frac{\partial IB(\tilde{\varepsilon}_0, \tilde{y}_0)}{\partial \tilde{y}_0} = L\rho(1-\delta)(\varepsilon - \tilde{\varepsilon}_0) \tag{B.7}$$

So, with a weak condition on the parameters, that $\varepsilon < \rho$, we again get monotonicity with respect to beliefs for all $\tilde{y}_0$[1]. Again banning is less attractive further from indifference. Lemma 6 said that at $\tilde{y}_0 = 0$, with beliefs above the critical value, it is optimal to ban. Then B.7 implies, that for any $\tilde{\varepsilon}_0 < \varepsilon$ it must be optimal to ban at all $\tilde{y}_0$. For $\tilde{\varepsilon}_0 > \varepsilon$ the incentive to ban is weaker at higher $\tilde{y}_0$ and so there may be a set of $\tilde{y}_0$ where it is optimal to allow if $\tilde{\varepsilon}_0 > \mu_H(\tilde{y}_0) (> \varepsilon)$.

The boundary traced out by $\mu_H(\tilde{y}_0)$ must be downward sloping. Let the value

[1] I will assume that since I am considering problems with small probabilities of bad events, the condition $\varepsilon < \rho$ holds. The proof goes through even if it doesn't, there will just be a set of low initial policy records for which the incentive to ban increases with $\tilde{\varepsilon}_0$, but since at the indifference line it always is strictly positive this just means that there is no allowing. In any case, if there is a set of states where allowing is optimal in this region, it must lie at higher $\tilde{\varepsilon}_0$.

of $\tilde{y}_0$ at which $\mu_H(\tilde{y}_0) = \frac{\hat{\varepsilon}}{\delta}$ be denoted γ. This will be useful in our discussion of the region $\tilde{\varepsilon}_0 > \frac{\hat{\varepsilon}}{\delta}$.

The thresholds in this region are given by,

$$\mu_{int}(\tilde{y}_0, \tilde{\varepsilon}_0) = \hat{\varepsilon} + \frac{(1-\delta)\left[(\rho\tilde{y}_0 + 1 - \rho)(\varepsilon - \hat{\varepsilon}) + (1-\varepsilon)\hat{\varepsilon}\right]}{(1-\delta)\rho\tilde{y}_0 + \delta(\rho - \varepsilon)} \tag{B.8}$$

The region with $\tilde{\varepsilon}_0 \in [\frac{\hat{\varepsilon}}{\delta}, 1]$

In the high $\tilde{\varepsilon}_0$ region, where even after observing no disaster the public still support a ban, we have

$$\begin{aligned} IB(\tilde{\varepsilon}_0, \tilde{y}_0) =& \rho\tilde{y}_0(a - \tilde{\varepsilon}_0 L) - (1-\varepsilon)(\rho\tilde{y}_0 + 1 - \rho)(a - \delta\tilde{\varepsilon}_0 L) \\ &- \varepsilon(\rho\tilde{y}_0 + 1 - \rho)(a - (\delta\tilde{\varepsilon}_0 + 1 - \delta)L) \end{aligned} \tag{B.9}$$

Now the partial derivatives are,

$$\frac{\partial IB(\tilde{\varepsilon}_0, \tilde{y}_0)}{\partial \tilde{\varepsilon}_0} = L\left[\delta(1-\rho) - \rho\tilde{y}_0(1-\delta)\right] \tag{B.10}$$

$$\frac{\partial IB(\tilde{\varepsilon}_0, \tilde{y}_0)}{\partial \tilde{y}_0} = L\rho(1-\delta)(\varepsilon - \tilde{\varepsilon}_0) \tag{B.11}$$

The thresholds in this region are given by,

$$\mu_{high}(\tilde{y}_0, \tilde{\varepsilon}_0) = \hat{\varepsilon} + \frac{(1-\delta)(\rho\tilde{y}_0 + 1 - \rho)(\varepsilon - \hat{\varepsilon})}{\rho\tilde{y}_0(1-\delta) - \delta(1-\rho)}$$

For $\tilde{y}_0 < \frac{(1-\rho)\delta}{\rho(1-\delta)}$, the derivative with respect to beliefs is positive, and for $\tilde{y}_0 > \frac{(1-\rho)\delta}{\rho(1-\delta)}$ it is negative. The incentive to ban decreases with $\tilde{y}_0$ iff $\tilde{\varepsilon}_0 > \varepsilon$, so again given Lemma 6, if there is any allowing it will be for the highest policy records.

Describing the full boundary of the regions A_H and B is slightly more complicated. Define γ such that $\mu_{int}(\gamma) = \frac{\hat{\varepsilon}}{\delta}$. If $\gamma < 1$ then, we can define a threshold $\mu_H(\tilde{y}_0)$ which is equal to $\mu_{int}(\tilde{y}_0)$ for $\tilde{y}_0 \in [\gamma, 1]$. The A_H region lies to the right of this boundary. As we said the boundary traced out by $\mu_{int}(\tilde{y}_0)$ will be downward sloping.

The incentive to ban is continuous crossing $\tilde{\varepsilon}_0 = \frac{\hat{\varepsilon}}{\delta}$ so at this value of the belief $\mu_{int}(\gamma) = \mu_{high}(\gamma)$. To describe what happens for higher initial beliefs we need to distinguish two cases. (a) Suppose $\gamma > \frac{(1-\rho)\delta}{\rho(1-\delta)}$. By equation B.10 for all $\tilde{y}_0 > \frac{(1-\rho)\delta}{\rho(1-\delta)}$, the incentive to ban continues to decrease eastwards of $\frac{\hat{\varepsilon}}{\delta}$ and allowing continues to be optimal. At $\tilde{y}_0 = \gamma$ the boundary is now given by $\mu_H(\tilde{y}_0) = \mu_{high}(\tilde{y}_0)$ which will continue to be downwards sloping and hit $\mu_{high}(\tilde{y}_0) = 1$ for some $\tilde{y}_0 > \frac{(1-\rho)\delta}{\rho(1-\delta)}$. (b) If $\gamma < \frac{(1-\rho)\delta}{\rho(1-\delta)}$, then for the set of $\tilde{y}_0 \in \left(\gamma, \frac{(1-\rho)\delta}{\rho(1-\delta)}\right)$, the incentive to ban increases as beliefs increase for $\tilde{\varepsilon}_0 > \frac{\hat{\varepsilon}}{\delta}$, and therefore it may become optimal to start banning again far enough from indifference. For $\tilde{y}_0 \in \left(\gamma, \frac{(1-\rho)\delta}{\rho(1-\delta)}\right)$, we have a second threshold then, $\mu_{HH}(\tilde{y}_0) = \mu_{high}(\tilde{y}_0)$, and this portion of the boundary will be upwards sloping.

The two cases are separated by the following condition: if $\varepsilon > \hat{\varepsilon}$, then $\gamma > \frac{(1-\rho)\delta}{\rho(1-\delta)}$ and we are in case (a), where the boundary between regions is downwards sloping everywhere. If $\varepsilon < \hat{\varepsilon}$ we are in the second case where the boundary slopes down until $\tilde{y}_0 = \gamma$ and the slopes up. Finally if $\varepsilon = \hat{\varepsilon}$, then at $\tilde{y}_0 = \gamma$ the boundary becomes horizontal.

Finally note that at $(\tilde{\varepsilon}_0 \geq \hat{\varepsilon}, 0)$, there is a strict incentive to ban, and that $\frac{\partial IB(\tilde{\varepsilon}_0,\tilde{y}_0)}{\partial \tilde{y}_0}$ is constant for fixed $\tilde{\varepsilon}_0$; so either it is positive, and so for those $\tilde{\varepsilon}_0$ there will be no allowing, or it is positive which means, if it is optimal to allow for some $\tilde{y}_0$ it must be optimal for $\tilde{y}_0 = 1$ so region A_H (if it exists) always contains some states on the northern edge of the state space.

B.2 Proof of Proposition 4

Effect of the volatility of the policy record

In the region of $\tilde{\varepsilon}_0 \leq \hat{\varepsilon}$,

$$\frac{\partial \mu_L(\tilde{y}_0)}{\partial \rho} = (1-\delta) \frac{\left[(1-\tilde{y}_0)(\varepsilon - \hat{\varepsilon}) - \varepsilon \tilde{y}_0 (1-\hat{\varepsilon}) - \delta \varepsilon (1-\tilde{y}_0)(1-\varepsilon)\right]}{[\ldots]^2} \tag{B.12}$$

If $\varepsilon \leq \left[\sqrt{(1-\delta)^2 + 4\delta\hat{\varepsilon}} - (1-\delta)\right]/2\delta$, then this is negative for all $\tilde{y}_0$. If ε is too high however, there will be some γ such that $\frac{\partial \mu_L(\tilde{y}_0)}{\partial \rho} > 0$ for $\tilde{y}_0 < \gamma$ and $\frac{\partial \mu_L(\tilde{y}_0)}{\partial \rho} < 0$ for $\tilde{y}_0 > \gamma$.

In the region of $\tilde{\varepsilon}_0 \geq \hat{\varepsilon}$, we have

$$\frac{\partial \mu_{int}(\tilde{y}_0)}{\partial \rho} = (1-\delta) \frac{\left[\varepsilon\delta(1-\tilde{y}_0)(\varepsilon - 1) - \varepsilon \tilde{y}_0 (1-\hat{\varepsilon})\right]}{[\ldots]^2} < 0$$

$$\frac{\partial \mu_{high}(\tilde{y}_0)}{\partial \rho} = (1-\delta) \frac{\left[-\tilde{y}_0(\varepsilon - \hat{\varepsilon})\right]}{[\ldots]^2}$$

The comparative static for μ_{high} is is positive (negative) for $\varepsilon < \hat{\varepsilon}$ $(\varepsilon > \hat{\varepsilon})$ but this is precisely the value at which the boundary of A_H changes from having an upwards sloping section to being every downwards sloping. Therefore, μ_H is decreasing with ρ and μ_{HH} is increasing in ρ: the set A_2 always expands when ρ is higher.

Effect of the volatility of the beliefs

In the region of $\tilde{\varepsilon}_0 \leq \hat{\varepsilon}$, we have

$$\frac{\partial \mu_L(\tilde{y}_0)}{\partial \delta} = \frac{\left[\hat{\varepsilon}(\varepsilon - \rho(1-\tilde{y}_0)) - \varepsilon(\rho\tilde{y}_0 + 1 - \rho)\right]}{[\ldots]^2}(-1-\varepsilon+\rho) > 0$$

In the region of $\tilde{\varepsilon}_0 \geq \hat{\varepsilon}$, we have

$$\frac{\partial \mu_{int}\left(\tilde{y}_0\right)}{\partial \delta} = \frac{\left(\varepsilon - \rho\right)\left[\left(\rho \tilde{y}_0 + 1 - \rho\right)\left(\varepsilon - \hat{\varepsilon}\right) + \left(1 - \varepsilon\right)\hat{\varepsilon}\right]}{\left[\ldots\right]^2} < 0$$

$$\frac{\partial \mu_{high}\left(\tilde{y}_0\right)}{\partial \delta} = \frac{\left(\rho \tilde{y}_0 + 1 - \rho\right)\left(1 - \rho\right)\left(\varepsilon - \hat{\varepsilon}\right)}{\left[\ldots\right]^2}$$

The comparative static for μ_{high} is negative (positive) for $\varepsilon < \hat{\varepsilon}$ ($\varepsilon > \hat{\varepsilon}$). Therefore, when $\varepsilon < \hat{\varepsilon}$, μ_H is decreasing with δ and μ_{HH} is decreasing with δ. For $\varepsilon > \hat{\varepsilon}$ still have $\frac{\partial \mu_H(\tilde{y}_0)}{\partial \delta} < 0$ for $\tilde{y}_0 > \gamma$, $\frac{\partial \mu_H(\tilde{y}_0)}{\partial \delta} > 0$ for $\tilde{y}_0 < \gamma$. In summary, in the region close to the indifference line, higher δ always leads to more allowing, but far from indifference there is more banning.

B.3 Proof of Proposition 5

The thresholds in the last period are:

$$\mu_{L,2}\left(0\right) = \hat{\varepsilon} - \left(1 - \delta\right)\frac{\left[\left(\rho - \varepsilon\right)\hat{\varepsilon} + \varepsilon\left(1 - \rho\right)\right]}{1 - \delta\left(\rho - \varepsilon\right)}$$

$$\mu_{L,2}\left(1\right) = \hat{\varepsilon} - \left(1 - \delta\right)\frac{\varepsilon\left(1 - \hat{\varepsilon}\right)}{1 - \rho + \delta\varepsilon}$$

$$\mu_{H,2}\left(1\right) = \begin{cases} \hat{\varepsilon} + \left(1 - \delta\right)\frac{\varepsilon(1-\hat{\varepsilon})}{\rho - \delta\varepsilon} & \text{if } \varepsilon < \frac{\rho\varepsilon}{\delta} \\ \hat{\varepsilon} + \left(1 - \delta\right)\frac{(\hat{\varepsilon} - \varepsilon)}{\rho - \delta} & \text{if } \varepsilon < \frac{\rho\varepsilon}{\delta} \text{ and } \rho > \delta \end{cases}$$

$$\mu_{HH,2}(1) = \widehat{\varepsilon} + (1 - \delta)\frac{\widehat{\varepsilon} - \varepsilon}{\delta - \rho} \text{ if } \rho < \delta \text{ and } \widehat{\varepsilon} > \varepsilon$$

In period 1, I need to be careful about distinguishing between all the different cases for how the DM will act in the second period for every first period action-event pair. First, suppose that the second period threshold for $\tilde{y}_0 = 0$ is more conservative than the first period one. Then there are four cases for behaviour at each of the two nodes that may be reached by allowing (one for each realisation of the event z). I calculate the expected reputational utility over two periods for each action in period 1 taking as given period 2 behaviour at each node. There are also different cases depending on whether if there is disaster initially and then no disaster in period 2, beliefs can be moderated below the indifference line, I will assume not: $\delta(1-\delta) > \widehat{\varepsilon}$. The proof for the cases where it is, is almost identical.

(i) Ban after ban, allow after allow (whether disaster or no disaster). Then the period 1 threshold is given by

$$\begin{aligned} \mu_{L,1} =& \frac{\widehat{\varepsilon}(1 + \beta + (1 + \beta)\varepsilon(1 - \varepsilon) + \varepsilon + \beta\varepsilon^2)}{1 + \beta + (1 + \beta)\delta^2\varepsilon(1 - \varepsilon) + \varepsilon\delta + \beta\varepsilon^2\delta^2} \\ & - (1 - \delta)\frac{[[\varepsilon(2 - \varepsilon) + \beta(\varepsilon^2(1 - \delta) + \varepsilon(1 - \varepsilon)\delta)]}{1 + \beta + (1 + \beta)\delta^2\varepsilon(1 - \varepsilon) + \varepsilon\delta + \beta\varepsilon^2\delta^2} \end{aligned} \tag{B.13}$$

(ii) Ban after ban, ban after allow (whether disaster or no disaster). Then the period 1 threshold is given by

$$\mu_{L,1} = \widehat{\varepsilon} - (1-\delta)\frac{[\widehat{\varepsilon}(1-\varepsilon)\beta(1-\rho) + (1-\widehat{\varepsilon})\varepsilon(1+\beta\rho)]}{1+\beta+\beta(\rho-1)\delta+\varepsilon\delta(1+\beta)} \tag{B.14}$$

(iii) Ban after ban, allow after allow and no disaster, ban after allow and disaster. Then the period 1 threshold is given by

$$\mu_{L,1} = \widehat{\varepsilon} - (1-\delta)\frac{[\varepsilon(1+\beta(1-\varepsilon+\rho)) - \widehat{\varepsilon}(\beta(1-\varepsilon)\varepsilon(1+\delta)+\varepsilon(1+\beta\rho))]}{1+\beta+\beta\varepsilon(1-\varepsilon)+\varepsilon(1+\beta\rho)} \tag{B.15}$$

iv) Ban after ban, ban after allow and no disaster, allow after allow and disaster. Then the period 1 threshold is given by

$$\mu_{L,1} = \widehat{\varepsilon} - (1-\delta)\frac{[\varepsilon(1+\beta\delta+\beta\rho) - \widehat{\varepsilon}((1-\varepsilon)(\rho-1)+\varepsilon(1+\beta+\beta\delta))]}{1+\beta+(1-\varepsilon)+(1-\varepsilon)(\rho-1)+\varepsilon(1+\beta)} \tag{B.16}$$

Next I consider the cases where the threshold $\mu_{L,2}(0)$ is less conservative than that in period 1. In this case, banning initially will be followed by allowing.

v) Allow after ban, allow after allow (whether disaster or no disaster). Then the period 1 threshold is given by

$$\mu_{L,1} = \widehat{\varepsilon} + (1-\delta)\frac{[\widehat{\varepsilon}\,(\beta(\rho-\varepsilon)+\varepsilon+\beta\varepsilon(1+\delta)) - \varepsilon\,(1+\beta\varepsilon(1-\rho)+\beta\varepsilon(1+\delta))]}{1+\beta\delta(\varepsilon-\rho)+\beta\varepsilon\delta^2+\varepsilon\delta} \tag{B.17}$$

vi) Allow after ban, ban after allow (whether disaster or no disaster). Then the period 1 threshold is given by

$$\mu_{L,1} = \widehat{\varepsilon} - (1-\delta)\frac{[\varepsilon(1+\beta) - \widehat{\varepsilon}(\varepsilon+\beta(2\rho-1))]}{1+\beta\delta(2\rho-1)+\varepsilon\delta} \tag{B.18}$$

vii) Allow after ban, allow after allow and no disaster, ban after allow and disaster. Then the period 1 threshold is given by

$$\begin{aligned}\mu_{L,1} = \widehat{\varepsilon} &- (1-\delta)\frac{\varepsilon\,(1+\beta(1-\varepsilon)+\beta\rho-\beta\varepsilon(1-\rho))}{1+\beta\rho\delta-\beta\varepsilon\delta+\beta\varepsilon(1-\varepsilon)\delta^2+\varepsilon\delta+\beta\varepsilon\rho\delta} \\ &- (1-\delta)\frac{\widehat{\varepsilon}\,(\beta(\rho-\varepsilon)+\varepsilon(1+\beta\rho)+\beta\varepsilon(1-\varepsilon)(1+\delta))}{1+\beta\rho\delta-\beta\varepsilon\delta+\beta\varepsilon(1-\varepsilon)\delta^2+\varepsilon\delta+\beta\varepsilon\rho\delta}\end{aligned} \tag{B.19}$$

viii) Allow after ban, ban after allow and no disaster, allow after allow and disaster. Then the period 1 threshold is given by

$$\mu_{L,1} = \widehat{\varepsilon} - (1-\delta)\frac{[\varepsilon\,(1+\beta\delta+\varepsilon\beta-\beta(1-\rho)) - \widehat{\varepsilon}\,(\beta(1-\varepsilon)\rho-\beta(1-\rho)+\varepsilon_\beta\varepsilon(1+\delta))]}{1+\beta(1-\varepsilon)\rho\delta-\beta(1-\rho)\delta+\varepsilon\delta+\beta\varepsilon\delta^2} \tag{B.20}$$

Note that the incentive to ban is always increasing in $\tilde{\varepsilon}_0$ (the denominator of the expressions is equal to $\frac{1}{L}\frac{\partial IB}{\partial \tilde{\varepsilon}_0}$ and is always positive).

B.4 Proof of Proposition 6

We assume $\delta = 0$ so after a z observation beliefs are extreme, $\tilde{\varepsilon}_t \in \{0, 1\}$. In the region with $\tilde{\varepsilon}_0 \leq \hat{\varepsilon}$, the value of banning is

$$\begin{aligned} V(0) &= (a - \tilde{\varepsilon}_0 L) \Big\{ (\rho \tilde{y}_0 - 1) + \beta(\rho^2 \tilde{y}_0 - 1) - \beta^2((\rho^3 \tilde{y}_0 - 1) \ldots \Big\} \\ &= (a - \tilde{\varepsilon}_0 L) \left[\frac{\rho \tilde{y}_0}{1 - \beta\rho} - \frac{1}{1 - \beta} \right] \end{aligned} \tag{B.21}$$

and the value of allowing is

$$\begin{aligned} V(1) =& (1 - \varepsilon)\rho(\tilde{y}_0 - 1)a + \varepsilon(\rho(\tilde{y}_0 - 1) + 1)(a - L) \\ &+ \beta \Big\{ (1 - \varepsilon)\rho^2(\tilde{y}_0 - 1)a + \varepsilon(\rho^2(\tilde{y}_0 - 1) + 1)(a - L) \Big\} \\ &+ \beta^2 \Big\{ (1 - \varepsilon)\rho^3(\tilde{y}_0 - 1)a + \varepsilon(\rho^3(\tilde{y}_0 - 1) + 1)(a - L) \Big\} \ldots \\ =& \frac{\rho(\tilde{y}_0 - 1)(a - \varepsilon L)}{1 - \beta\rho} - \frac{\varepsilon(a - L)}{1 - \beta} \end{aligned} \tag{B.22}$$

It is straightforward to show that the incentive to ban $IB(\tilde{y}_0, \tilde{\varepsilon}_0)$ is increasing with $\tilde{\varepsilon}_0$ and that there exists a ξ such that the incentive to ban increases with $\tilde{y}_0$ if $\varepsilon > \xi$ and is decreasing otherwise. So we have a threshold solution, and the threshold will decrease with $\tilde{y}_0$ if risk is high enough.

$$\mu_L(\tilde{y}_0) = \hat{\varepsilon} - \frac{\varepsilon(L - a)(1 - \beta\rho) + \rho(1 - \tilde{y}_0)(a - \varepsilon L)(1 - \beta)}{1 - \beta\rho - \rho\tilde{y}_0(1 - \beta)} \tag{B.23}$$

It must be optimal to ban along the indifference line and also at an initial record of zero. The incentive to ban $IB(\tilde{y}_0, \tilde{\varepsilon}_0)$ is now decreasing with $\tilde{\varepsilon}_0$, so if there is any allowing it must be for high enough $\tilde{\varepsilon}_0$ for a given $\tilde{y}_0$. Further, since the boundary of the allowing region cannot meet the indifference line or the bottom edge of the state space, if there exists an area of allowing, it will be bounded by a downward sloping boundary. So we have a threshold solution, and the threshold will decrease with $\tilde{y}_0$.

$$\mu_H(\tilde{y}_0) = \hat{\varepsilon} + \frac{\rho(1 - \tilde{y}_0)(a - \varepsilon L)(1 - \beta) + \varepsilon(L - a)(1 - \beta\rho)}{\rho\tilde{y}_0(1 - \beta)} \tag{B.24}$$

Appendix C

Chapter 3 - Proofs and Extensions

C.1 Proof of Proposition 7

At the optimal policy, the value functions must satisfy the following functional equations (after some simplification),

$$\begin{aligned} V_L(L) &= \alpha^C p_L\left(x_L^L, x_R^L, L\right)\left(x_R^L - x_L^L\right) - \alpha^C\left(x_R^L - \widehat{x}^L\right) \\ &\quad +\beta\left[p_L\left(x_L^L, x_R^L, L\right)\left(V_L(L) - V_L(R)\right) + V_L(R)\right] \end{aligned} \tag{C.1}$$

$$\begin{aligned} V_L(R) = &\quad \alpha^C p_L\left(x_L^R, x_R^R, R\right)\left(x_R^R - x_L^R\right) - \alpha^C\left(x_R^R - \widehat{x}^L\right) \\ &\quad +\beta\left[p_L\left(x_L^R, x_R^R, R\right)\left(V_L(L) - V_L(R)\right) + V_L(R)\right] \end{aligned} \tag{C.2}$$

$$\begin{aligned} V_R(L) = &\quad \alpha^C p_L\left(x_L^L, x_R^L, L\right)\left(x_L^L - x_R^L\right) - \alpha^C\left(\widehat{x}^R - x_R^L\right) \\ &\quad +\beta\left[p_L\left(x_L^L, x_R^L, L\right)\left(V_R(L) - V_R(R)\right) + V_R(R)\right] \end{aligned} \tag{C.3}$$

$$\begin{aligned} V_R(R) &= \alpha^C p_L\left(x_L^R, x_R^R, R\right)\left(x_L^R - x_R^R\right) - \alpha^C\left(\widehat{x}^R - x_R^R\right) \\ &\quad +\beta\left[p_L\left(x_L^R, x_R^R, R\right)\left(V_R(L) - V_R(R)\right) + V_R(R)\right] \end{aligned} \tag{C.4}$$

where

$$\begin{aligned} p_L\left(x_L^L, x_R^L, L\right) &= \frac{1}{2} + \widehat{b}\phi + 2\alpha^v\phi\left(\overline{x}^L - \widehat{x}^m\right) \\ p_L\left(x_L^R, x_R^R, R\right) &= \frac{1}{2} - \widehat{b}\phi + 2\alpha^v\phi\left(\overline{x}^R - \widehat{x}^m\right) \end{aligned}$$

and $\overline{x}^J = \frac{\left(x_L^J + x_R^J\right)}{2}$.

I focus on the symmetric equilibrium, where

$$x_L^L = -x_R^R \text{ and } x_L^R = -x_R^L \tag{C.5}$$

This implies that the degree of divergence is independent of the state: $\delta x^L = \delta x^R = \delta x$, and for the bias we have $\overline{x}^L = -\overline{x}^R$. Inspecting the expressions for the probabilities of L winning, this also implies $p_L\left(x_L^L, x_R^L, L\right) = 1 - p_L\left(x_L^R, x_R^R, R\right)$.

I derive a useful expression for solving the problem, by subtracting the two value functions for each player,

$$V_L(L) - V_L(R) = \frac{\alpha^C\left(p_L(L) - p_L(R)\right)\delta x + \alpha^C\left(x_R^R - x_R^L\right)}{1 - \beta\left(p_L(L) - p_L(R)\right)} \tag{C.6}$$

where I use the shorthand notation $p_L(L) = p_L\left(x_L^L, x_R^L, L\right)$ etc.

Taking first order conditions, starting with the one for x_L^L,

$$\alpha^C p_L(L) = \phi\alpha^v\left[\alpha^C\left(x_R^L - x_L^L\right) + \beta\left(V_L(L) - V_L(R)\right)\right] \tag{C.7}$$

and for x_L^R,

$$\alpha^C p_L(R) = \phi\alpha^v\left[\alpha^C\left(x_R^R - x_L^R\right) + \beta\left(V_L(L) - V_L(R)\right)\right] \tag{C.8}$$

These are the Euler equations for the problem. The left hand side of the equation shows, for each state, the marginal cost to L of moving his platform towards the median, conditional on winning he will get a lower utility. The right hand side shows the marginal benefit, his probability of winning increases, and this changes the expected policy utility today, and he may additionally benefit in future, if the continuation value is higher in the state where he is incumbent, compared to the state where he is challenger. Note that his decision about x_L depends on the expected marginal value of winning (where marginal refers to the difference between winning and losing)[1].

Next, using all the symmetry results, namely $\delta x = \left(x_R^L - x_L^L\right) = \left(x_R^R - x_L^R\right)$ and $p_L(R) = 1 - p_L(L)$, the two expressions above imply

$$p_L(L) = p_L(R) = 1 - p_L(L) = \frac{1}{2}$$

Again I obtain the result, as in the one-period model, that parties adjust their platforms in such a way that the probability of winning is always equal for both parties, no matter who is the incumbent.

From the result for the probabilities of winning I can show that the mean policy, in each state, is the same as I found in the one period game:

$$\overline{x}^L = -\frac{\widehat{b}}{2\alpha^v} \text{ and } \overline{x}^R = \frac{\widehat{b}}{2\alpha^v} \tag{C.9}$$

Also from either Euler equation, using our result for equilibrium values of p_L, and using expression (C.6) for the difference in value functions in the two states, I get

$$\frac{1}{2}\alpha^C = \phi\alpha^v\alpha^C\left(\delta x - 2\beta\overline{x}^L\right)$$

where the $\overline{x}^L$ term comes from $\left(x_R^R - x_R^L\right) = -x_L^L - x_R^L = 2\overline{x}^L$. I can now solve for the degree of divergence,

$$\delta x = \frac{1}{2\phi\alpha^v} - \frac{\beta\widehat{b}}{\alpha^v} \tag{C.10}$$

C.2 Proof of Proposition 8

I can compute expressions for the expected quality of candidates. Challengers, for which only one signal is ever observed, which has accuracy q, have the following posteriors.

$$\tilde{\mu}_{i,t}\left(s_{i,t} = 1\right) = q + (1-q)\,\mu \tag{C.11}$$

$$\tilde{\mu}_{i,t}\left(s_{i,t} = 0\right) = \mu\,(1-q) \tag{C.12}$$

This applies to both candidates in period 1 and to the single challenger in period 2. For the incumbent in period 2, both signals will be taken into account, considering that their accuracies may differ.

[1] This distinction will become more important in section 3.4 when I incorporate the learning process into the model of ideological policy choice.

$$\tilde{\mu}_{i,t}\left(s_{i,1}=1, s_{i,2}=1\right)=\frac{\left[q+(1-q)\,\mu\right]\left[p+(1-p)\,\mu\right]}{\mu+pq\,(1-\mu)} \tag{C.13}$$

$$\tilde{\mu}_{i,t}\left(s_{i,1}=1, s_{i,2}=0\right)=\frac{\left[q+(1-q)\,\mu\right](1-p)}{1-pq} \tag{C.14}$$

$$\tilde{\mu}_{i,t}\left(s_{i,1}=0, s_{i,2}=1\right)=\frac{(1-q)\left[p+(1-p)\,\mu\right]}{1-pq} \tag{C.15}$$

$$\tilde{\mu}_{i,t}\left(s_{i,1}=0, s_{i,2}=0\right)=\frac{\mu\,(1-q)\,(1-\mu)\,(1-p)}{1-\mu+\mu pq} \tag{C.16}$$

It is useful to consider the ranking of posteriors in period $t=2$. Without loss of generality assume L won the first election. I will write these posteriors as $\tilde{\mu}_{R,2}\left(s_{R,2}\right)$ for R, since the R candidate from the first period has been replaced, so $s_{R,1}$ is not used in the expectations of challenger quality in this period, and $\tilde{\mu}_{L,2}\left(s_{L,1}, s_{L,2}\right)$ for L.

For any values of the parameters, it holds that

$$\tilde{\mu}_{L,2}\,(0,0)<\tilde{\mu}_{R,2}\,(0)<\mu<\tilde{\mu}_{R,2}\,(1)<\tilde{\mu}_{L,2}\,(1,1)$$

as well as

$$\tilde{\mu}_{L,2}\,(0,0)<\tilde{\mu}_{L,2}\,(1,0)\lesseqgtr\tilde{\mu}_{L,2}\,(0,1)<\tilde{\mu}_{L,2}\,(1,1)$$

There are three cases for the complete ranking of posteriors. In case 1, I have high p so that the comparison of the incumbent's posterior to that of the challenger depends solely on the incumbent's second period signal. In cases 2 and 3 I have that mixed signals give moderate posteriors, so that the ranking of incumbent and challenger now does depend on the challenger's signal, at least for one of the challenger's signals.

Case I: $s_{L,2}$ is decisive

This is the case where p is sufficiently large (relative to q) for $s_{L,2}$ to determine the ranking of candidates in period $t=2$:

$$\tilde{\mu}_{L,2}\,(0,0)<\tilde{\mu}_{L,2}\,(1,0)<\tilde{\mu}_{R,2}\,(0)<\tilde{\mu}_{R,2}\,(1)<\tilde{\mu}_{L,2}\,(0,1)<\tilde{\mu}_{L,2}\,(1,1)$$

and is satisfied for all μ, p, and q such that:

$$p>\frac{q}{\mu+q\,(1-2\mu+\mu q)} \tag{C.17}$$

With this ordering of the posteriors, the incumbent's second period signal is decisive for the election, he wins if and only if $s_{L,2}=1$. Therefore the probability of winning the second election is

$$Pr\,(L\,wins\,at\,t=2\mid L\,wins\,at\,t=1)=Pr\,(s_{L,2}=1\mid L\,wins\,at\,t=1)$$

$$=Pr\,(v_L=1\mid L\,wins\,at\,t=1)\,[p+(1-p)\,\mu]+Pr\,(v_L=0\mid L\,wins\,at\,t=1)\,[(1-p)\,\mu]$$

$$= (1-p)\,\mu + Pr\,(v_L = 1 \mid L\,wins\,at\,t = 1)\,p$$

Although winning the first election does improve the conditional quality of the incumbent ($Pr\,(v = 1 \mid L\,wins\,at\,t = 1) > \mu$), if the first period signal was very weak it does not do so substantially so that the expression above remains below $\frac{1}{2}$. The full expression for the probability of high quality conditional in winning an election is

$$Pr\,(v_L = 1 \mid L\,wins\,at\,t = 1) = \frac{\mu\left[(q + (1-q)\,\mu)\left(\frac{\mu}{2} + 1 - \mu\right) + (1-q)\,(1-\mu)\,\frac{(1-\mu)}{2}\right]}{\frac{1}{2}}$$

so that, after simplification

$$\Delta = pq\mu\,(1-\mu) + \mu - \frac{1}{2}$$

This expression is continuously and strictly increasing in q. As $\mu \to 0$, this tends to $\mu - \frac{1}{2}$ which is negative, since the prior is unbalanced in favour of low quality. Therefore, there must exist a threshold $q^*\,(p,\mu) \in (0,p]$ such that $\hat{b} < 0$ for $q \leq q^*\,(p,\mu)$. Specifically, for the values of μ, p, q satisfying (C.17), I have

$$q_I^*\,(\mu, p) = \frac{\frac{1}{2} - \mu}{p\mu\,(1-\mu)} \tag{C.18}$$

Case II: $s_{L,2} = 0$ **is not decisive but** $s_{L,2} = 1$ **is**

This is the case where the ordering is

$$\tilde{\mu}_{L,2}\,(0,0) < \tilde{\mu}_{R,2}\,(0) < \tilde{\mu}_{L,2}\,(1,0) < \tilde{\mu}_{R,2}\,(1) < \tilde{\mu}_{L,2}\,(0,1) < \tilde{\mu}_{L,2}\,(1,1)$$

In this case if $s_{L,2} = 1$, L wins the second election; but he might win even with $s_{L,2} = 0$, as long as $s_{R,2} = 0$ and $s_{L,1} = 1$. This is the most favourable case for the incumbent, in terms of the set of outcomes under which he gets re-elected. The condition for this case is

$$\frac{q}{q\mu + (1-\mu)\,(q^2 - q + 1)} < p < \frac{q}{\mu + q\,(1 - 2\mu + \mu q)} \tag{C.19}$$

The probability of L winnning re-election is

$$Pr\,(L\,wins\,at\,t = 2 \mid L\,wins\,at\,t = 1) = Pr\,(s_{L,2} = 1 \mid L\,wins\,at\,t = 1)$$

$$+Pr\,(s_{R,2} = 0) \cdot Pr\,(s_{L,1} = 1, s_{L,2} = 0 \mid L\,wins\,at\,t = 1)$$

After computing the different expressions I have

$$\Delta = \mu\left[1 + (1-\mu)^2\,(2-\mu)\right] + qp\mu\,(1-\mu)\,[1 - (1-\mu)\,(2-\mu)] - \frac{1}{2}$$

This expression is increasing in q for $\mu \geq \hat{\mu} \approx 0.38$. Further, for $q = 0$, there is incumbency disadvantage as long as, $\mu \leq \underline{\mu} \approx 0.25$. Therefore, for $\mu \leq \underline{\mu}$, case II predicts incumbency disadvantage for all q, and for $\mu \geq \hat{\mu}$ it predicts incumbency advantage for all q. The final case is $\mu \in \left(\underline{\mu}, \hat{\mu}\right)$, $\hat{b}$ may cross zero from above, at the value

$$q_{II}^{**} = \frac{\frac{1}{2} - \mu \left[1 + (1-\mu)^2 (2-\mu)\right]}{p\mu (1-\mu) \left[1 + (1-\mu)(2-\mu)\right]}$$

Inspecting case III will allow us to say something more about this case.

Case III: $s_{L,2}$ is not decisive

This is the case where the ordering is either

$$\tilde{\mu}_{L,2}(0,0) < \tilde{\mu}_{R,2}(0) < \tilde{\mu}_{L,2}(1,0) < \tilde{\mu}_{L,2}(0,1) < \tilde{\mu}_{R,2}(1) < \tilde{\mu}_{L,2}(1,1)$$

or,

$$\tilde{\mu}_{L,2}(0,0) < \tilde{\mu}_{R,2}(0) < \tilde{\mu}_{L,2}(0,1) < \tilde{\mu}_{L,2}(1,0) < \tilde{\mu}_{R,2}(1) < \tilde{\mu}_{L,2}(1,1)$$

In any case, mixed signals are intermediate, and will result in posteriors that lie between those of the challenger.

In this case, when $s_{R,2} = 0$ the incumbent wins as long as as he doesn't have two bad signals, and when $s_{R,2} = 1$ the incumbent wins only if he has a record of performing well in both periods.

$$Pr\left(L\, wins\, at\, t = 2 \mid L\, wins\, at\, t = 1\right) = \mu Pr\left(s_{L,1} = s_{L,2} = 1 \mid (L\, wins\, at\, t = 1\right)$$

$$+ (1-\mu)\left[1 - Pr\left(s_{L,1} = s_{L,2} = 0 \mid (L\, wins\, at\, t = 1)\right]\right.$$

The posteriors that I need to compute for the incumbent are, therefore,

$$Pr\left(s_{L,1} = s_{L,2} = 1 \mid L\, wins\, at\, t = 1\right)$$

$$= Pr\left(s_{L,1} = 1 \mid L\, wins\, at\, t = 1\right) \sum_{v} Pr\left(v \mid s_{L,1} = 1\right) Pr\left(s_{L,2} = 1 \mid v\right)$$

$$Pr\left(s_{L,1} = s_{L,2} = 0 \mid L\, wins\, at\, t = 1\right)$$

$$= Pr\left(s_{L,1} = 0 \mid L\, wins\, at\, t = 1\right) \sum_{v} Pr\left(v \mid s_{L,1} = 0\right) Pr\left(s_{L,2} = 0 \mid v\right)$$

Expanding out for these posteriors, I end up with

$$Pr\left(L\, wins\, at\, t = 2 \mid L\, wins\, at\, t = 1\right)$$

$$= 2\mu^2 \left(1 - \frac{\mu}{2}\right) \left[(q + (1-q)\mu)(p + (1-p)\mu) + \mu(1-\mu)(1-q)(1-p)\right]$$

$$+ (1-\mu)\left[1 - (1-\mu)^2 \left[\mu(1-\mu)(1-q)(1-p) + (q + (1-q)(1-\mu)(p + (1-p)(1-\mu))\right]\right]$$

With some simplification of the above I obtain the expression for the incumbency advantage

$$\Delta = 2\mu^2\left(1-\frac{\mu}{2}\right)\left[\left(1-\mu\right)qp+\mu\right]\left(1-\mu\right)\left\{1-\left(1-\mu\right)^2\left[\mu qp+1-\mu\right]\right\}-\frac{1}{2} \quad \text{(C.20)}$$

Now, firstly notice that, rather counterintuitively, this expression is not always increasing in q despite the fact that higher q improves selection. I have

$$\frac{\partial\Delta}{\partial q} = \mu\left(1-\mu\right)p\left[\mu\left(2-\mu\right)-\left(1-\mu\right)^2\right], \quad \text{(C.21)}$$

which is positive if and only if $\mu > \bar{\mu} = 1 - \frac{\sqrt{2}}{2}$. So for very low μ it is less likely that the incumbent will win, the higher q. The reason for this is that higher q increases the correlation of the signals of any given player, so both the probability that $s_{L,1} = s_{L,2} = 1$ and the probability that $s_{L,1} = s_{L,2} = 0$ for the incumbent go up. When μ is very low, this shift of probability from mixed records to matched signals hurts the incumbent, as the increase in probability of $s_{L,1} = s_{L,2} = 0$ outweighs the effect from higher probability of $s_{L,1} = s_{L,2} = 1$.

I can write the value of q for which $\Delta = 0$,

$$q^*_{III} = \frac{\frac{1}{2} - 2\mu^3\left(1-\frac{\mu}{2}\right) - 1 + \mu + \left(1-\mu\right)^4}{\mu\left(1-\mu\right)p\left[\mu\left(2-\mu\right)-\left(1-\mu\right)^2\right]} \quad \text{(C.22)}$$

Doing the necessary algebra, one can show that the numerator of this expression is positive precisely for $\mu < \bar{\mu}$ when the denominator is negative, and viceversa for $\mu > \bar{\mu}$; which implies that $q^* < 0$ always for this case. Therefore, it is immediate that, for all p, q that fall into case III, $\Delta \geq 0$ when $\mu \geq \bar{\mu}$, and that $\Delta < 0$ when $\mu < \bar{\mu}$.

Further, I can show two things. The first, is that $\mu < \bar{\mu}$ implies that the threshold in expression (C.18) is above 1, and so for $\mu < \bar{\mu}$, and as long as $p > 0$, and $q < 1$ (i.e. as long as the voters learn something more about the incumbent than the challenger), case I also predicts an incumbency disadvantage. The second is that, at the value of q at which we switch from case II to case III, there is a discontinuous drop in Δ because case III drops one of the outcomes under which the incumbent used to win in case II. This implies that for $\mu \geq \bar{\mu}$, case II predicts incumbency advantage[2].

There is one special case, where $\mu \in \left(\underline{\mu}, \bar{\mu}\right)$, where cases I and III predict incumbency disadvantage, but case II may contain a set of values of q for which there is incumbency advantage. This is the only case in which the sign of Δ is not monotonic in q[3].

Figure C.1 shows an example, fixing $p = 0.5$, and evaluated at three different values of the prior on quality, μ.

[2]At the highest value of q for which I am in case II, Δ is strictly higher than the value predicted by case III, which is positive. And since $\hat{\mu} > \bar{\mu}$, case II has Δ decreasing in q so that for lower q the incumbency advantage is even greater.

[3]Note that the value of Δ is never monotonic in q because it must go up as switch from case I to II and down when switch, at higher q, from case II to III.

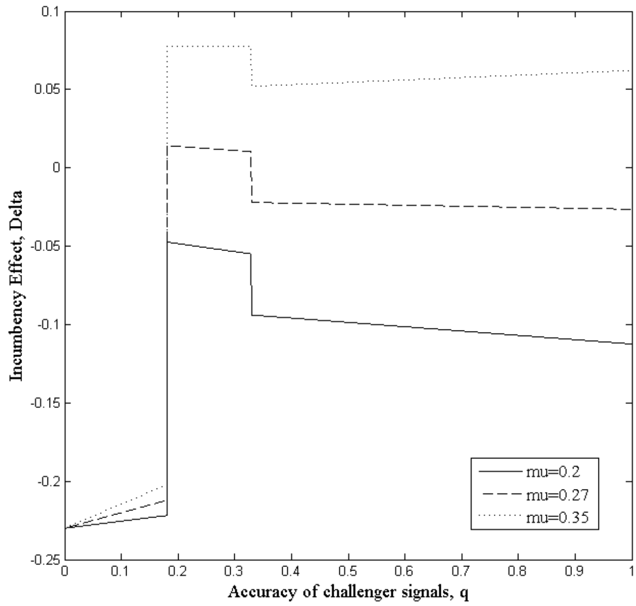

Figure C.1: Incumbency effect for different values of the prior on quality, μ

Bibliography

Acemoglu, Daron and James A Robinson (2000). "Why did the West extend the franchise? Democracy, inequality, and growth in historical perspective". In: *Quarterly Journal of Economics*, pp. 1167–1199.

— (2001). "A theory of political transitions". In: *American Economic Review*, pp. 938–963.

Achen, Christopher H and Larry M Bartels (2004). "Musical Chairs: Pocketbook Voting and the Limits of Democratic Accountability." Presented at the Annual Meeting of the American Political Science Association". In:

— (2012). "Blind Retrospection: Why Shark Attacks Are Bad For Democracy". In: *Center for the Study of Democratic Institutions, Vanderbilt University. Working Paper.*

Aghion, Philippe and Patrick Bolton (1997). "A theory of trickle-down growth and development". In: *The Review of Economic Studies* 64.2, pp. 151–172.

Alesina, Alberto and Roberto Perotti (1994). "The political economy of growth: a critical survey of the recent literature". In: *The World Bank Economic Review* 8.3, pp. 351–371.

Alesina, Alberto and Dani Rodrik (1991). *Distributive politics and economic growth.* Tech. rep. National Bureau of Economic Research.

— (1994). "Distribution politics and income distribution". In: *Quarterly Journal of Economics* 109, pp. 465–490.

Alesina, Alberto and Guido Tabellini (1990). "A positive theory of fiscal deficits and government debt". In: *The Review of Economic Studies* 57.3, pp. 403–414.

Alford, John R and David W Brady (1989). "Personal and partisan advantage in US congressional elections". In: *Congress Reconsidered, 4th edition, edited by Lawrence C. Dodd and Bruce I. Oppenhemier. New York: Praeger Publishers.*

Ansolabehere, Stephen, David Brady, and Morris Fiorina (1992). "The vanishing marginals and electoral responsiveness". In: *British Journal of Political Science* 22.01, pp. 21–38.

Ansolabehere, Stephen and James M Snyder (2002). "The incumbency advantage in US elections: An analysis of state and federal offices, 1942-2000". In: *Election Law Journal* 1.3, pp. 315–338.

Arcelus, Francisco and Allan H Meltzer (1975). "The effect of aggregate economic variables on congressional elections". In: *American Political Science Review* 69.04, pp. 1232–1239.

Ashworth, Scott and Ethan Bueno de Mesquita (2008). "Electoral selection, strategic challenger entry, and the incumbency advantage". In: *The Journal of Politics* 70.04, pp. 1006–1025.

Austen-Smith, David (2000). “Redistributing income under proportional representation”. In: *Journal of Political Economy* 108.6, pp. 1235–1269.
Azzimonti, Marina (2014). “The dynamics of public investment under persistent electoral advantage”. In: *Review of Economic Dynamics*.
Banks, Jeffrey S and John Duggan (2005). “Probabilistic voting in the spatial model of elections: The theory of office-motivated candidates”. In: *Social choice and strategic decisions*. Springer, pp. 15–56.
Bartels, Larry M (2009). *Unequal democracy: The political economy of the new gilded age*. Princeton University Press.
Baumeister, Roy F, Ellen Bratslavsky, Catrin Finkenauer, and Kathleen D Vohs (2001). “Bad is stronger than good.” In: *Review of general psychology* 5.4, p. 323.
Benabou, Roland (1996). “Inequality and growth”. In: *NBER Macroeconomics Annual 1996, Volume 11*. MIT Press, pp. 11–92.
— (2000). “Unequal societies: Income distribution and the social contract”. In: *American Economic Review*, pp. 96–129.
Bertola, Giuseppe (1991). *Factor shares and savings in endogenous growth*. Tech. rep. National Bureau of Economic Research.
Besley, Timothy and Stephen Coate (1998). “Sources of inefficiency in a representative democracy: a dynamic analysis”. In: *American Economic Review*, pp. 139–156.
Bloom, Howard S and H Douglas Price (1975). “Voter response to short-run economic conditions: The asymmetric effect of prosperity and recession”. In: *American Political Science Review* 69.04, pp. 1240–1254.
Breaux, David (1990). “Specifying the Impact of Incumbency On State Legislative Elections A District-Level Analysis”. In: *American Politics Research* 18.3, pp. 270–286.
Canes-Wrone, Brandice, Michael C Herron, and Kenneth W Shotts (2001). “Leadership and pandering: A theory of executive policymaking”. In: *American Journal of Political Science*, pp. 532–550.
Coughlin, Peter J (1992). *Probabilistic voting theory*. Cambridge University Press.
Cox, Gary W (1990). “Centripetal and centrifugal incentives in electoral systems”. In: *American Journal of Political Science*, pp. 903–935.
Cox, Gary W and Jonathan N Katz (1996). “Why did the incumbency advantage in US House elections grow?” In: *American Journal of Political Science*, pp. 478–497.
Cox, Gary W and Scott Morgenstern (1993). “The increasing advantage of incumbency in the US states”. In: *Legislative Studies Quarterly*, pp. 495–514.
— (1995). “The incumbency advantage in multimember districts: Evidence from the US States”. In: *Legislative Studies Quarterly*, pp. 329–349.
Dewan, Torun and Rafael Hortala-Vallve (2013). “Electoral competition, control and learning”. In: *Unpublished manuscript*.
Feddersen, Timothy and Faruk Gul (2014). “Polarization and Income Inequality: A Dynamic Model of Unequal Democracy”. In:
Ferejohn, John A (1977). “On the decline of competition in congressional elections”. In: *American Political Science Review* 71.01, pp. 166–176.
Fiorina, Morris P (1978). “Economic retrospective voting in American national elections: A micro-analysis”. In: *American Journal of Political Science*, pp. 426–443.
— (1981). “Retrospective voting in American national elections”. In:

Fischhoff, Baruch and Ruth Beyth (1975). "I knew it would happen: Remembered probabilities of once—future things". In: *Organizational Behavior and Human Performance* 13.1, pp. 1–16.
Fudenberg, Drew and Jean Tirole (1991). *Game Theory.* Translated into Chinesse by Renin University Press, Bejing: China. Cambridge, MA: MIT Press.
Galor, Oded and Joseph Zeira (1993). "Income distribution and macroeconomics". In: *The review of economic studies* 60.1, pp. 35–52.
Gelman, Andrew and Gary King (1990). "Estimating incumbency advantage without bias". In: *American Journal of Political Science*, pp. 1142–1164.
Healy, Andrew J, Neil Malhotra, and Cecilia Hyunjung Mo (2010). "Irrelevant events affect voters' evaluations of government performance". In: *Proceedings of the National Academy of Sciences* 107.29, pp. 12804–12809.
Hinich, Melvin J (1977). "Equilibrium in spatial voting: The median voter result is an artifact". In: *Journal of Economic Theory* 16.2, pp. 208–219.
Hoban, Thomas J (1997). "Consumer acceptance of biotechnology: an international perspective." In: *Nature biotechnology* 15.3, p. 232.
Huber, Gregory A, Seth J Hill, and Gabriel S Lenz (2012). "Sources of bias in retrospective decision making: Experimental evidence on voters' limitations in controlling incumbents". In: *American Political Science Review* 106.04, pp. 720–741.
Jacobson, Gary C (1987). "The marginals never vanished: Incumbency and competition in elections to the US House of Representatives, 1952-82". In: *American Journal of Political Science*, pp. 126–141.
Jewell, Malcolm E and David Breaux (1988). "The effect of incumbency on state legislative elections". In: *Legislative Studies Quarterly*, pp. 495–514.
Kahneman, Daniel (2011). *Thinking, fast and slow.* Macmillan.
Kahneman, Daniel and Amos Tversky (1973). "On the psychology of prediction." In: *Psychological review* 80.4, p. 237.
— (1979). "Prospect theory: An analysis of decision under risk". In: *Econometrica: Journal of the Econometric Society*, pp. 263–291.
Klašnja, Marko and Rocıo Titiunik (2013). "Incumbency Disadvantage In Weak Party Systems: Evidence from Brazil". In: *Unpublished manuscript, University of Michigan.*
Kramer, Gerald H (1971). "Short-term fluctuations in US voting behavior, 1896–1964". In: *American Political Science Review* 65.01, pp. 131–143.
Krehbiel, Keith and John R Wright (1983). "The incumbency effect in congressional elections: A test of two explanations". In: *American Journal of Political Science*, pp. 140–157.
Levitt, Steven D and Catherine D Wolfram (1997). "Decomposing the sources of incumbency advantage in the US House". In: *Legislative Studies Quarterly*, pp. 45–60.
Lindbeck, Assar and Jörgen W Weibull (1993). "A model of political equilibrium in a representative democracy". In: *Journal of Public Economics* 51.2, pp. 195–209.
Linden, Leigh L (2004). "Are incumbents really advantaged? The preference for non-incumbents in Indian national elections". In: *Unpublished paper.*
Maskin, Eric and Jean Tirole (2004). "The politician and the judge: Accountability in government". In: *American Economic Review*, pp. 1034–1054.

McCarty, Nolan, Keith T Poole, and Howard Rosenthal (2006). *Polarized America: The dance of ideology and unequal riches*. Vol. 5. mit Press.

Meltzer, Allan H and Scott F Richard (1981). "A rational theory of the size of government". In: *The Journal of Political Economy*, pp. 914–927.

Molina, José (2001). "The electoral effect of underdevelopment: government turnover and its causes in Latin–American, Caribbean and industrialized countries". In: *Electoral Studies* 20.3, pp. 427–446.

Nunnari, Salvatore and Jan Zápal (2014). "Dynamic Elections and Ideological Polarization". In:

Pauly, Mark V (2007). "Risks and benefits in health care: the view from economics". In: *Health Affairs* 26.3, pp. 653–662.

Perotti, Roberto (1992). "Income distribution, politics, and growth". In: *The American Economic Review*, pp. 311–316.

— (1993). "Political equilibrium, income distribution, and growth". In: *The Review of Economic Studies* 60.4, pp. 755–776.

— (1996). "Growth, income distribution, and democracy: what the data say". In: *Journal of Economic growth* 1.2, pp. 149–187.

Persson, Torsten and Lars EO Svensson (1989). "Why a stubborn conservative would run a deficit: Policy with time-inconsistent preferences". In: *The Quarterly Journal of Economics*, pp. 325–345.

Persson, Torsten and Guido Tabellini (1991). *Is inequality harmful for growth? Theory and evidence*. Tech. rep. National Bureau of Economic Research.

— (1992). "Growth, distribution and politics". In: *European Economic Review* 36.2, pp. 593–602.

— (1994). "Representative democracy and capital taxation". In: *Journal of Public Economics* 55.1, pp. 53–70.

Popkin, Samuel L (1994). *The reasoning voter: Communication and persuasion in presidential campaigns*. University of Chicago Press.

Prato, Carlo (2011). "Electoral Competition and Partisan Policy Feedbacks". In: *Available at SSRN 2027499*.

Saint-Paul, Gilles and Thierry Verdier (1993). "Education, democracy and growth". In: *Journal of Development Economics* 42.2, pp. 399–407.

Slovic, Paul (1987). "Perception of risk". In: *Science* 236.4799, pp. 280–285.

Slovic, Paul, Baruch Fischhoff, and Sarah Lichtenstein (1979). "Rating the risks". In: *Environment: Science and Policy for Sustainable Development* 21.3, pp. 14–39.

Titiunik, Rocıo (2009). "Incumbency advantage in Brazil: Evidence from municipal mayor elections". In: *University of California-Berkeley, mimeo*.

Tufte, Edward R (1980). *Political control of the economy*. Princeton University Press.

Tversky, Amos and Daniel Kahneman (1971). "Belief in the law of small numbers." In: *Psychological bulletin* 76.2, p. 105.

— (1973). "Availability: A heuristic for judging frequency and probability". In: *Cognitive psychology* 5.2, pp. 207–232.

— (1974). "Judgment under uncertainty: Heuristics and biases". In: *science* 185.4157, pp. 1124–1131.

Uppal, Yogesh (2009). "The disadvantaged incumbents: estimating incumbency effects in Indian state legislatures". In: *Public Choice* 138.1-2, pp. 9–27.

Zaller, John (1998). "Politicians as Prize Fighters: Electoral Selection and Incumbency Advantage". In: *Politicians and party politics*, pp. 125–85.