Political Science and Political Economy Working Paper

Department of Government

London School of Economics

No. 4/2012

Teams of Rivals: Learning about a Cabinet and its Shadow

Torun Dewan (LSE)

Rafael Hortala-Vallve (LSE)

THE LONDON SCHOOL OF ECONOMICS AND POLITICAL SCIENCE

Teams of Rivals: Learning about a Cabinet and its Shadow*

Torun Dewan

London School of Economics and Political Science

t.dewan@lse.ac.uk

Rafael Hortala-Vallve

London School of Economics and Political Science

r.hortala-vallve@lse.ac.uk

Abstract. We explore electoral accountability in a model in which an incumbent team chooses between a "safe" option with a known payoff and a "reform" policy that yields higher or lower payoffs depending on whether the politician implementing it is competent or not. The first best outcome allows for learning about politicians' type, but is subject to the free-rider problem: if team members are moderately competent, on average, they prefer others to implement the reforms. A cabinet provides an institutional remedy by allowing politicians to obtain individual rents from office and making policy choice subject to unanimous consent. Although free-riding is alleviated, outcomes with a cabinet are qualitatively similar to those in its absence. We relate this to a hold-up problem: politicians revealed as incompetent veto the implementation of reform. Surprisingly, the problem is most severe when such politicians can be removed from office. Electoral competition between rival teams, as in that between a cabinet and a shadow cabinet, restores learning; but at moderately high levels of competence, relative to the efficient benchmark, such competition leads to over-investment in risky reforms. This contrasts with the case where executive decision-making is by an individual–then electoral competition leads to first best outcomes.

*For helpful comments we thank Gilat Levy, Hande Mutlu-Eren, participants at the "Priorat Workshop on Formal Models in Political Science 2012", and audience members at the Annual Meetings of the European Political Science Association, Berlin, 2012.

1. Introduction

In many situations of political interest rival teams of politicians compete for executive office. Most formal models of political competition, however, focus on individual decision makers: Downs described parties as "teams of politicians", but in this canonical model of electoral competition the parties are unitary actors; the principal-agent literature that studies executive accountability has few well developed models of competition, and even then the focus is on individual agents. As a consequence our understanding is limited to cases where competition is between rival leaders (as in presidential elections or primaries) or candidates (as in constituency run-offs); whilst competition between rival teams, such as that between an incumbent cabinet and its shadow, has been ignored.

A critical aspect of decision-making in teams that has been commented upon in the political science literature, is the advantage faced by a principal who learns about the competencies of a set of agents who perform similar tasks. Huber and Martinez-Gallardo (2003, 2008) study executive turnover in parliamentary democracies and relate empirical patterns in the data to such a learning process:

> "Since party leaders will often be uncertain as to which politicians have the technical expertise and skills necessary to do their jobs well, a process of trial and error occurs to discover the best talent. This can only occur by getting rid of some ministers, bringing in new faces, and reshuffling individuals from one post to another."

Similarly, Berlinski, Dewan, and Dowding (2010) show that patterns in the data on ministerial turnover in British government from 1945-2007 are related to policy failures that "provide information to the Prime Minister that was not available when she appointed her minister." Using the framework developed by Holstrom (1982), they show that the data is consistent with what we would expect to observe were a principal (the Prime Minister) to use a retention rule that reflects comparative evaluation of cabinet ministers' performance on correlated tasks

These empirical studies suggest that acquisition of information about politicians' competencies is central to accountability in cabinets. However, it remains to be shown that cabinet government is an effective and efficient means for learning about such competencies when compared with other executive forms. We provide a theoretical framework that allows us to explore the cabinet performance against an efficient level of learning. We contrast the performance of an executive team (with respect to learning) to one that is appointed, explore the role played by competitive elections between rival (potential) executive teams, and compare outcomes under cabinet government with those that occur when decisions are taken by an individual.

In our model an agent is characterized by his degree of competence–either high or low–which we define as the ability to implement a policy that *ex-post* is in the public's best interest. The agent's type is unknown, even to himself, and is revealed only ex-post when implementing a particular type of policy. Specifically we characterize some policies as being "safe" in that the outcome is known, whereas a "reform" policy yields a high payoff only when implemented by a competent agent.

As a benchmark we analyze an executive team that consists of agents who care only about the policy outcome. The first best case illustrates perfectly the logic of learning: the executive team implements the risky policy for a wide range of the parameter space even when their individual expected first period payoffs are less than those when taking the safe option. Doing so allows the team to learn the individual competencies of its members. By this trial and error process, those who have the competence necessary to do their jobs maintain their positions, whilst others do not. Indeed we find that there is more learning in larger teams. Focussing on what is individually rational, rather than what is optimal for the team, however, we find that learning is lower than in the first best case and is decreasing in the size of the team. The problem is the classic one of free-riding: an individual agent does not internalize the positive effect of his action on the team.

Can the institutional framework of a cabinet alleviate this problem? Two aspects of cabinet governance suggest that it should. Firstly, cabinets are (typically) formed by politicians who obtain rents from holding office and so these individual incentives should alleviate free-riding. Secondly, cabinets operate collective responsibility: the policy executed by an individual politician is chosen by the executive team under the "unanimity principle" that all government office-holders must support the same policy. This principle ensures that each politician is decisive in determining the policy to be implemented and so mitigates free-riding. We analyze a situation where prior to the execution of policy in each period, a cabinet of politicians must collectively agree whether the reform policy should be implemented.

Despite the absence of free-riding, we find that learning is lower than in the first best case and is decreasing in the size of the team. In fact our results are qualitatively similar to those obtained without a cabinet. We relate this finding to the classic hold-up problem. A politician who is revealed as incompetent will vote against the further implementation of the reform policy. Anticipating this, the first period value of investment in the risky policy falls for *all* politicians, the Cabinet is less likely to take that course of action, and this effect grows larger in the size of the Cabinet. A hold-up occurs because a politician does not realize the full return on his investment due to the bargaining power of other politicians who can veto policy.

Of course there exists a resolution to this problem. Under the doctrine of individual ministerial responsibility the party, or its leader, can replace a politician who is revealed as incompetent. This should alleviate the hold-up problem since it takes away the bargaining power of any individual politician. Surprisingly, however, we show that it destroys *all* learning.

Whilst our basic model casts doubt on the efficacy of the cabinet as a means of acquiring relevant information about politicians' competencies, it abstracts from elective and competitive elements of cabinet governance. According to both Dicey and Mill (in *Representative Government*), the normative force of collective responsibility is derived not from what is efficient from the cabinets's perspective but from the will of the majority. According to Marshall (1984), the interpretation of these scholars was that the convention be designed so as to "secure that the will of the majority of the electorate be put into effect." Moreover, Mill and Bagehot, defended representative government as a process by which men and their ideas could be tested (against custom) in a *competitive* political process: this they called "government by discussion".

These competitive and elective elements are central to cabinet governance. Whilst Parliament is in session, the daily dual between its front-benches provides the adversarial drama of British politics: on the front-benches on one side of the Commons sits the Cabinet and other government ministers; on the opposite side of the House sits Her Majesties loyal opposition, with the front-bench seats reserved for members of its Shadow Cabinet. Turner (1969) notes that "the Shadow Cabinet both derives from and influences the basic features of the British system of government, a two-party system *in which voters choose directly between rival teams of potential Ministers*."[1]

We build these competitive and elective aspects into our model in several stages.

First we analyze a model (a version of the canonical principal-agent model of accountability) where voters select between a cabinet and a prospective rival team (a shadow cabinet). When voters are constrained to either retain the cabinet or not then (we show that) their ex-ante preferred level of risk matches that which is optimal for the (average) cabinet member. Although competition between rival teams recovers learning, the level of risk-taking is either too high or too low relative to the efficient benchmark. When the Cabinet has a low average level of competence it is unwilling to take the risky action since then it reveals its incompetence and is replaced. By contrast at relatively high levels of competence the cabinet is too willing to implement reforms. Doing so allows it to "showcase" its competence and so be reelected (by a sequentially rational voter).

[1]Our emphasis added

Moving beyond the canonical model applied to teams, we next introduce an active role for the Shadow Cabinet. In particular we explore outcomes when the shadow can best respond given the incumbent cabinet's policy record. Specifically we allow the shadow cabinet to collectively choose whether to campaign on the reform policy or the safe one. And we suppose that if they choose the former, the competence of each of its members is revealed with some probability. Our justification for this modeling assumption is that voters form an assessment of a rival team's competence within the context of an electoral campaign and contrast this with their assessment of the incumbent executive that is based on its record. Although levels of risk-taking are different our earlier intuitions continue to hold. Competition between an active cabinet and its shadow does not recover efficiency.

Finally we study the voters optimal retention scheme. We show that in addition to selecting between rival teams of politicians, voters can provide incentives that limit inefficiencies. However, as a consequence of competitive tensions between the cabinet and shadow cabinet, the former is still too willing to implement reforms. A particularly interesting finding concerns the range of the parameter space for which we find free-riding in the absence of a cabinet, and showcasing in its presence. Both effects are apparent when the team is of high average levels of competence: free-riding entails learning that is too low; showcasing delivers a degree of learning that is excessive.

Our study yields three central results that are surprising. First, although we set up our model to explore whether there is an association between cabinet governance and learning about politicians' competencies, we find that although cabinet government alleviates free-riding it introduces another, equally severe, hold-up problem that mitigates learning. Second, the mechanism that links cabinet government to learning about politicians competencies is competitive interaction between rival teams of prospective office holders. This is surprising because the competitive interaction between a cabinet and its shadow is a subject largely ignored in the political science literature. Although our results confirm that it is desirable to include competitive team aspects in models of accountability, our third central result highlights negative consequences of such competition between teams. At high levels of competence the cabinet introduces an excess of reform policies that allows it to "showcase" its talent. Although voters can control such behavior via an optimal retention scheme in an election between a single incumbent and a challenger, they are unable to do so when electoral competition is between a cabinet and a shadow cabinet.

In our concluding comments we discuss the relevance of our findings for different organizational forms.

Our paper is structured as follows. We first discuss our contribution to the related literature. In section 3 we explore policy implementation by a team that is not structured as a cabinet. In section 4 we explore our policy environment when the decision is made by a cabinet. In section 5 we explore the same environment but introduce a role for a party or leader who is able to remove members of the cabinet team. In section 6 we study the competitive interaction between a cabinet and shadow cabinet where the incumbent cabinet can be replaced by election. In section 7 we explore the optimal retention scheme of the voter for different executive sizes. Finally section 8 concludes.

2. Related Literature

We contribute to a large literature on electoral accountability, going back to classic models by Barro (1973) and Ferejohn (1986) that were developed further by Banks and Sundaram (1993) and Banks and Sundaram (1998), and on the allocation of policies to elected or appointed officials Alesssina and Tabellini (2007a,b). In contrast to those papers our focus is on electoral accountability where decisions are made by teams of politicians and where voters select between an incumbent team evaluated on its performance and a rival one that campaigns on policy. Our focus on elections with an active opponent relates to Ashworth and Shotts (2011, 2010). In Ashworth and Shotts (2011) the action of the challenger is to provide an evaluation of the incumbent's performance rather than, as in our model, to propose a policy choice.[2] The focus in that paper is on individual politicians, whereas ours is on teams.

Our analysis of team decision-making in cabinets relates to a literature on accountability in parliamentary democracies that builds on the framework introduced by Strøm (1985). Typically these models have focussed on issues of moral hazard (Dewan and Myatt, 2007; Indridason and Kam, 2008; Dewan and Myatt, 2010). Dewan and Hortalla-Valve (2011) analyze selection of ministers but do not embed cabinet government in a model of elections.

We contribute to a small literature on the role of cabinets. Cox (2011) develops a moral-hazard argument that traces the historical emergence of cabinet government in 18th century Britain. Dewan, Galeotti, Ghiglino, and Squintani (2011) justify cabinet governance from the perspective of information aggregation. Here we develop a model that explores the role of a cabinet and its shadow with respect to learning. In so doing we build on the empirical analysis of Huber and Martinez-Gallardo (2008) and Berlinski, Dewan, and Dowding (2010, 2012) who informally relate

[2]An incumbent engages in costly information acquisition before choosing policy whilst a challenger can assess and criticizes the incumbent's platform. When such claims are verifiable, voters can use retrospective voting strategies to sanction incumbents and so induce incentives for better performance.

ministerial turnover to notions of team learning (Holstrom, 1982) and yardstick competition (Besley and Case, 1995).

Related papers on learning include Volden, Ting, and Carpenter (2008) who analyze a situation where government's learn from their own experiences and those of other governments, Callander (2011a) who analyzes learning by trial and error in a business environment, and Callander (2011b) who looks at learning in a spatial model of elections. A closely related paper is by Callander and Hardstad (2012) who analyze the free-rider problem in federal systems where policymakers can benefit from policy experimentation in neighboring states.

We explore learning within the context of the two-armed bandit model, used in policy analysis by Aghion, Bolton, Harris, and Julien (1991) and adapted by Strumpf (2002) to look at policy innovation and its relation to government decentralization. A recent contribution Hirsch (2011) analyzes learning where the principal and agent share the same intrinsic motivation but may differ with respect to their preferred policy instrument. A paper closely related to ours is Strulovici (2010) who analyzes experimentation by groups of decision-makers operating under different collective decision-making rules. Our focus is on electoral competition between teams.[3]

The specific policy choice we study, namely that between a safe and a risky option, has been studied by Lizzeri and Persico (2009) who assess the impact of different electoral systems on risk control. Papers by Majumdar and Mukand (2004) and Fu and Li (2010) look at the same choice in an electoral context, though both assume politicians know their type. The former highlight that reputation concerns may cause a politician to persist with risky choices that have failed. The closest model to ours is Dewan and Hortala-Vallve (2012) who used the multi-armed bandit model to show that competitive elections ameliorate electoral incentives to adopt risky policies (and gain reputation). Our analysis shows that whilst this result holds when electoral competition is between individuals it can be overturned when the contest is between teams.

The reason for this somewhat surprising finding is the existence of something akin to what, in the study of industrial organization, is referred to as a hold-up problem. The problem, first studied by Williamson (1992) as part of a broader class of problems with transaction-costs, is that the return on a specific investment by an individual firm is subject to the bargaining power of other firms. Whilst the literature on hold-up problems in economics is vast, we are aware of no direct applications in political science.

[3]Our paper also relates to a classic paper by Banks and Sundaram (1990) who analyze election an infinite armed bandit problem where a principal selects a candidate with a single action that yields a reward (to the principal) according to the agent's type.

3. Policy Experiments in Teams

We begin with a team of n agents who choose a policy in each of two periods. We (initially) assume that these agents care (only) about the final policies that are implemented by the team. Each agent faces a policy choice between a "safe" policy that yields a payoff of 1 to the agent who implements it regardless of her type, and a risky policy. The outcome from implementation of the risky policy can either be a success or a failure. If the agent who implements the risky policy is competent then a success is obtained and this yields a payoff of $r > 1$ to the agent. If however the agent who implements the risky policy is not competent then the outcome is failure. The probability any individual is competent is $p \in (0, 1)$. Each agent cares only about policy and so their payoff is the average policy payoff obtained by the team

The timing is as follows. First nature determines whether an agent is competent or not. Next each agent implement a policy, either safe or risky. If they choose the latter then the outcome, either success or failure, is commonly observed and, as a consequence, the competence of the agent who implements the policy is revealed. Agents can then choose whether to remain as part of the executive team or not. Upon doing so the remaining agents choose the second period policies and finally payoffs are realized.

In our first result we explore the role of the team in this environment.

Proposition 1. *If executing the first-best, a team of agents implement the risky policy in period 1 if and only if* $\frac{1}{r} < \frac{p+1-(1-p)^n}{2-(1-p)^n}$. *Implementation of the risky policy in the first period increases in the size of the team and more is learned in large teams than in small ones. When agents decide independently whether to implement the risky policy, the unique equilibrium has exactly k agents implement the risky policy in the first period where k is the maximum value $\tilde{k}$ that satisfies* $\frac{1}{r} < \frac{p+p(1-p)^{\tilde{k}-1}}{1+p(1-p)^{\tilde{k}-1}}$ *(when the maximum value is larger than n, we have that $k = n$; in other words, all members of the team implement the risky policy).*

The first part of Proposition 1 shows that in the first best scenario there is learning and this effect increases in the size of the team. This is illustrated in the left-hand panel of Figure 1. The dashed line on the 45 degree angle illustrates where the expected payoffs to playing reform and safe are equivalent. To the left of this line is the area where the expected first period payoffs from playing safe exceed those from implementing the reform.

First note that when the decision is made by an individual then the risky policy is chosen *even* when her expected first period payoff is less than that when taking the safe option. To the right

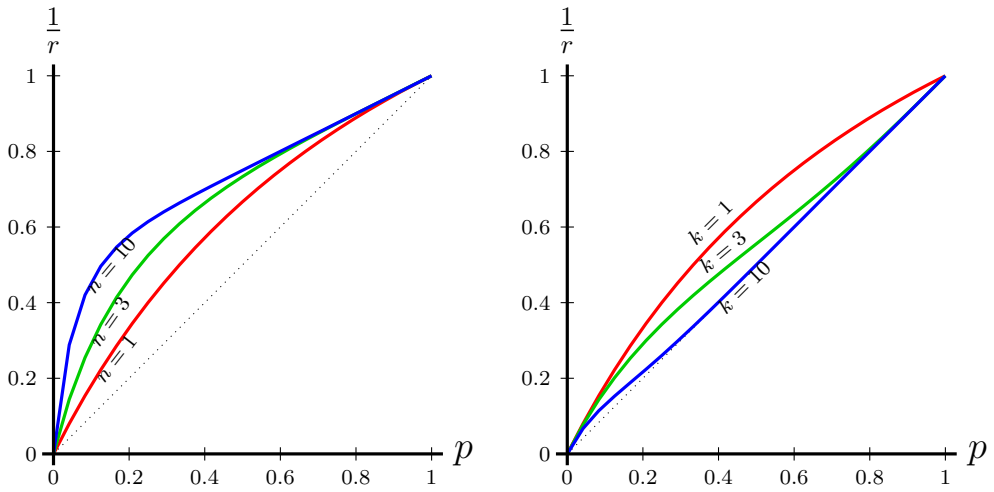

FIGURE 1. Policy Experimentation by a Team of Agents. The dotted line indicates $pr = 1$ where expected payoffs from both policies are the same. The left-hand panel shows the first best outcome. To the right of each curve, are the parameter configurations for which in the first-best the team experiments in period 1 for different values of n. The right-hand panel illustrates the individually rational outcomes: to the right of each curve, the parameter configurations for which k individuals experiment in period 1 (given that $n \geq k$).

of the line indicated $n = 1$ a single agent implements reform. Despite the risk in implementing the policy, the agent learns her type. If she is competent then this increases the second period payoff and so yields a higher payoff than playing safe.

Next we note, as illustrated by the curves $n = 3$ and $n = 10$ in the left-hand panel of Figure 1, that this learning effect increases in the size of the team making the decision. Given that the agents are perfectly aligned and care only about policy outcomes they are unaffected by the identity of the decision-maker; they care only about her type, whether competent or not. Indeed should a single agent be revealed as competent, then the team is better off when delegating all decisions to her. It follows that players' payoffs are increasing in the likelihood that a single member of the team is revealed to be competent. Moreover, as the team grows large the more likely it is that the implementation of the risky policy will yield at least one success. The main claim in Proposition 1 then follows: large teams encourage learning.

The second part of 1 focuses on individual incentives. Given the alignment in players' payoffs it is possible that each agent would wish another to experiment in the first period whilst herself choosing safe. And this is what we find. In short, as in Bolton and Harris (1999), team learning is subject to the free-rider problem. The effect is illustrated in the right-hand panel of Figure 1. Comparing the left and right hand panel we observe that, at high levels of competence in particular, whilst it is optimal for the entire team to experiment with reform in equilibrium only one agent does so.

In the following section we revise our model to allow for institutional features of team decision-making that are relevant to cabinet governance. We ask whether the institutional features of cabinet can mitigate the free-rider problem and thereby provide a rationale for cabinet governance.

4. Policy Experiments in the Cabinet

The first factor that distinguishes the Cabinet from our earlier example is that it is formed by ministers who are "career politicians" and obtain ministerial rents from office. In the United Kingdom, for example, the Cabinet is formed predominantly by members of the House of Commons, though sometimes includes members from the Lords. Whilst politicians care about policy outcomes this is not their only motive; they seek the prestige, benefits, and influence of high office and, upon attaining that goal, wish to remain there. Politicians seek office in part due to a desire for policy influence. We capture their twin (policy and office-seeking) concerns in a simple and parsimonious manner. We assume that a member of the Cabinet (a minister) obtains a payoff from the policy implemented $-r$ if the risky policy is implemented and she is competent, 0 if the risky policy is implemented and she is incompetent, and s if the safe policy is implemented– only when she is in office; otherwise the politician obtains 0. The key impact of these rents is in providing an individual incentive for a politician to be in office that curtails the desire to free-ride.

The definitive feature of Cabinet government is the convention of collective responsibility. Put crudely, this means (i) that ministers must all agree with government policy or resign their positions; and (ii) that the government as a whole is held responsible for its actions and no individual minister can abrogate himself from responsibility. Here we include (i) directly in our set up by assuming that the policies chosen in the first and second periods require the unanimous consent of Cabinet; later on in our model of competitive elections we include a version of (ii). Thus we include directly what Marshall refers to as the "unanimity principle". We note that the impact of collective team decision-making under unanimity is that it imposes the play of symmetric strategies: either the Cabinet collectively agree to play "safe" or to adopt the "risky" platform.

The timing is as follows. First nature determines whether a politician is competent or not. Next the cabinet jointly agrees in a single vote under unanimity whether each implements a safe or reform policy. If they choose the latter then the outcomes for each minister, whether a success or failure, is commonly observed and, as a consequence, the competence of the politician who implements the policy is revealed. Whatever policy is chosen in the first period the cabinet then jointly agrees

with a single vote under unanimity whether each implements a safe or reform policy in the second period. Outcomes are obtained and then final payoffs are realized.

In the context of a cabinet, our restriction to the binary choice in which an agent chooses between safe and reform requires some justification. We have in mind a situation where the policy represents a broad agenda of reform. The cabinet must agree collectively to this programme, which is then implemented by individual ministers who head specific departments and thus tailor individual policies to suit the broader agenda. One main example is the implementation of an austerity programme of budget cuts. Whilst the overall policy must be agreed at the cabinet level, individual ministers have jurisdiction over how and where such cuts are made. Successful implementation of such reforms therefore reveals a minister's competence.

We contrast decision-making by a cabinet with that of a team of agents discussed earlier, summarize our main findings in Proposition 2 below and provide an illustration in Figure 2.

Proposition 2. *A cabinet operating under collective responsibility implements the risky policy in period 1 only when* $pr > \frac{1+p^n}{1+p^{n-1}} \Leftrightarrow \frac{1}{r} < \frac{p+p^n}{1+p^n}$. *The willingness to implement the risky policy in the first period decreases in the size of the team. Less is learned in large cabinets than in small ones.*

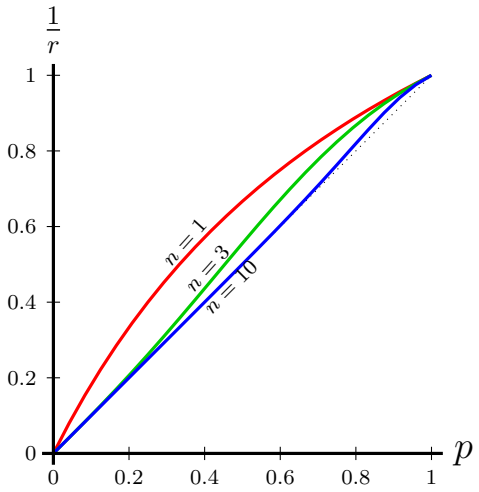

FIGURE 2. Policy Experimentation by a Cabinet. The dotted line indicates $pr = 1$ where expected payoffs from both policies are the same. To the right of each curve are parameter configurations for which teams of size n experiment in period 1.

Although the cabinet surmounts the free-rider problem, learning does not increase in the size of the cabinet team. In fact, and as immediate apparent from Figure 2, experimentation with reform is weakly decreasing in the size of the team. Moreover the effect is more pronounced at lower competence.

Note that when cabinet chooses the risky policy then some politicians may be revealed to be incompetent. For example, if we treat the risky policy as being the implementation of budget cuts or administrative reforms then, whilst some politicians will be able to make these adjustments successfully, (other) ministers will fail. Those who fail to successfully implement the reforms would then obtain a zero payoff were the same policies to be implemented for a further period. Thus it is optimal for those politicians to veto the implementation of reforms in the second period.

The intuition follows directly. Although a politician anticipates a benefit of learning his type, as before, when implementing the risky policy, this additional information does not necessarily relate to an increased payoff in the second period should he be revealed as competent. In fact, the benefits of learning are dependant upon *all* politicians successfully implementing the risky policy. Only then will the Cabinet implement the cuts or reforms for a further period and the successful minister receive an increased payoff.

The restriction in learning is thus due to the bargaining power of members of the cabinet. A minister who successfully implements the risky policy does not reap the full reward when others hold-up the further implementation of reforms. Such hold-up problems have long been studied in economics–the key idea is that in an arrangement between two firms, neither firm is willing to make optimal investments that benefits both if they thereby increase the bargaining power of the other–but have received little attention in political science.[4]

The hold-up problem is exacerbated in larger teams. This is in direct contrast to our earlier analysis of the first-best scenario where the larger the team the more likely it was that a single member would be *successful* and so all members benefit. Now the larger the team the more likely it is that a single member will be *unsuccessful* and thus trigger the hold-up. Indeed, as we observe in Figure 2, when the team has 10 or more members then learning is restricted to the case where $pr > 1$.

We note that these effects in our model are not due to the stark nature of the unanimity requirement. Indeed another way to look at the disadvantages of having a team of politicians is that each politician knows that, as the group size increases, he is less likely to be pivotal in tomorrow's decision. Then, as a consequence, he is less willing to experiment in the first period. And this is true under any collective decision-making rule.

[4]A somewhat related idea and application is found in Dewan and Myatt (2007). They show that ministers willingness to invest in reform policies decreases in the value of their political career. In their model reforms, whilst beneficial, can also expose ministers to attacks by interest groups.

5. Resignation in the Cabinet

The key obstacle to learning in this environment is the expectation that members of the Cabinet who have been unable to successfully implement risky policies, will veto their further implementation. A possible resolution to that problem is, of course, that such ministers be removed from the cabinet. In fact the convention of individual ministerial responsibility means that, although ministers are collectively responsible for policy, a minister cannot abrogate himself from responsibility for the execution of that policy. Thus a minister who is seen to have failed at implementing his task may be subject to a resignation call and subsequently be removed from the Cabinet.

We next explore a game where, after first period policy has been implemented and outcomes observed, a principal who maximizes the average gains of the Cabinet can adjust the rewards from office according to performance. We have in mind the actions taken by a party leader or prime minister who can fire a minister who has failed at his task. Although one might expect that the removal of the obstacle to learning, namely the bargaining power of ministers who have failed at implementing their task will lead to a restoration of learning, in fact the opposite is true as we record in the following proposition.

Proposition 3. *A cabinet with n members that operates under collective and individual responsibility that is imposed by a party or its leader never implements the risky policy in period* 1 *when* $pr < 1$ *and always implements it when* $pr > 1$.

Although learning is inhibited due to the hold-up problem, the removal of the bargaining power of obstructive ministers via their (forced) resignation from cabinet, exacerbates the issue. Indeed this destroys all learning. Recall that a minister is unaware of his type and learns about his competence only when implementing the risky policy. Whilst ex-post, and conditional on successful implementation of that policy, a minister would prefer his cabinet colleagues who fail be removed from government. However the very possibility of removal acts as a constraint upon investment in the risky option. A minister fears that when cabinet implements the risky policy, she will be revealed as incompetent and so lose her position. The prospect of resignation reduces the value obtained from the lottery when playing risky, relative to the sure payoff from playing safe. As a consequence each minister substitutes away from the risky policy and into the safe option.

6. The Cabinet-vs-Shadow Cabinet

As we have seen, in principle, team decision-making provides a rich environment for learning about the competencies of politicians. However, the specific conventions of cabinet governance, namely collective and individual ministerial responsibility, can provide a barrier to the implementation of policies that would allow for such learning. To conclude on this basis that cabinet governance is inimical to learning, however, would be premature.

Thus far we have considered the workings of a cabinet in isolation from the competitive elements of representative democracy; and so, as a consequence of our restricted focus, we may be missing an important aspect of team learning. As noted in our introductory remarks, the interplay between the cabinet and the shadow cabinet is central to the structure of democratic governance in the Westminster system. To analyze these competitive tensions we now develop our model to consider voter selection between rival teams of ministers: a cabinet and a shadow cabinet.

The key difference with our earlier model is that, after the outcomes associated with the first period policy implemented by the incumbent cabinet are observed, a shadow cabinet pronounces in favor of either the safe or the reform policy. Although the action set is the same for both the cabinet and its shadow there is a key difference: whilst the former stands on an observed record of implemented policy, and is judged by voters accordingly, the latter does not. Instead members of the shadow cabinet can reveal their competence only via their policy pronouncements. When the shadow cabinet advocates reform then it must deliver a justification and plan for implementation of the reform and this provides a test of its politicians' competence. By contrast, advocating the status quo is more straightforward and less testing of the shadow cabinet's abilities. We thus model the revelation of the shadow cabinet members competencies in the simplest possible way; we assume that when it pronounces in favor of the reform policy, the competence of each member of the shadow cabinet is revealed with probability $q_s \in [0, 1]$. By contrast, when the shadow cabinet advocates the safe policy then each member is competent with prior probability p.

Specifically, the timing of our game is as follows: First nature determines whether a politician, whichever team he is a member of, is competent or not (we assume that all politicians are competent with prior probability p). Then, and as before, each minister, a member of the incumbent cabinet, implements a policy that is agreed upon in Cabinet under the unanimity principle. Policy outcomes associated with the decisions of cabinet members are observed, before the shadow cabinet pronounces on policy. If the outcome of the implementation of reform is a success then a voter

learns that the cabinet member is competent. If, on the other hand, reform results in failure then the voter learns that the politician is competent. If the shadow cabinet chooses reform then the competence of each of its members is revealed with probability q_s. Voters then evaluate the average competence of both the cabinet and shadow cabinet before electing one or other in a winner-take-all contest. Whichever team forms the new government then implements the second period policy under unanimity.

In an equilibrium of our model, the policies implemented by the Cabinet, the policy proposals made by the Shadow Cabinet, and voters' final decisions are sequentially rational given beliefs held by the electorate about the competence of the Cabinet and its shadow on the equilibrium path of play.

First we assess the policy the voter would have the cabinet implement. This provides an efficient benchmark against which we can compare outcomes. Perhaps surprisingly we show that this conforms to the policy that would be implemented by a team of officials who are appointed to office.

Lemma 1. *In the presence of elections in which a cabinet seeks reelection, the voter wants experimentation in the same circumstances as when the voter appoints a team of officials (with no replacement) into government. The voter's optimal level of experimentation is such that the incumbent team should implement the risky action in period* 1 *only when* $pr > \frac{1+p^n}{1+p^{n-1}} \Leftrightarrow \frac{1}{r} < \frac{p+p^n}{1+p^n}$.

Since the voter must choose between rival teams who implement policy next we analyze her optimal behavior. When all cabinet members are competent, the voter should reelect the government with certainty as it ensures a payoff of $r > 1$ in period 2. What happens when there is at least one non-competent member? If $pr < 1$ the voter knows that the incumbent team or a new one implements the safe policy in period 2, and so is indifferent as to which team implements the second period policy. We suppose that the voter selects the team whose expected (average) competence is highest–this would indeed be the optimal action should there be more than two periods. Given this tie-breaking rule, the incumbent party always replaces non-competent team members. When the voter is still indifferent we assume the incumbent is reelected with probability 1/2.

Given this voting rule it is straightforward to show that should the cabinet implement the reform policy and any of its members be revealed as incompetent then it is always better for the party, or its leaders to replace such a minister. Therefore, in what follows, we assume that they do just that and so we explore the level of reforms implemented by an incumbent cabinet that faces re-election and operates under collective and individual responsibility imposed by a party or its leader.

To simplify the analysis we focus our analysis on the difference between q_s and 1 which captures a critical element in competition between rival teams. The cabinet stands on an observed policy record, by contrast the shadow cabinet stands on an untested policy pronouncement. It is without loss of generality to assume perfect correlation between the incumbent's type and the policy outcome if she chooses the risky policy. Doing so allows us to build our analysis on several interesting cases. In the first we set $q_s = 0$. Then, whereas voters evaluate the incumbent on her record, any pronouncement by the opposition shadow cabinet is dismissed by the electorate (i.e. the opponent's pronouncements are uninformative). In this institutional environment the retention of the cabinet is tied solely to its performance. This yields a standard model of elections in which they serve as a mechanism for selecting competent politicians. There is a key difference, however, with the standard model, namely that the voter anticipates policies that will be implemented by the cabinet on the basis of the unanimity principle. Given that the shadow cabinet is inactive when $q_s = 0$ as its pronouncements are ignored, we refer to this as a "noncompetitive election".

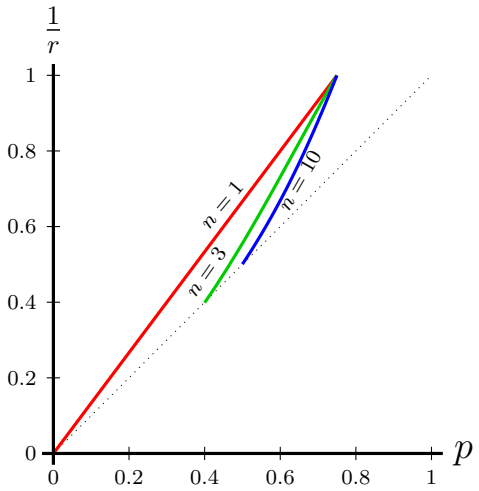

FIGURE 3. Policy Experimentation by a Cabinet that seeks reelection after period 1. The dotted line indicates $pr = 1$ where expected payoffs from both policies are the same. To the right of each curve, we have the parameter configurations for which the team experiments in period 1 for different values of n.

Proposition 4. *In a noncompetitive election, a cabinet with n members that operates under collective and individual responsibility that is imposed by a party or its leader, and which seeks reelection after period* 1*, implements the risky policy when* $pr < 1$ *and* $\frac{1}{r} < \frac{p+p^n}{1.5-p+p^n}$ *(and always when* $pr > 1$*). Relative to an appointed cabinet operating under collective responsibility, implementation of the reform policy in the first period is either too high (at high levels of competence) or too low (at low levels competence).*

The intuition for why there is under-investment in reform at low average levels of competence follows immediately from our earlier result in proposition 3. Since it is always optimal for the party or its leader to remove from office a minister who is revealed as incompetent, the cabinet will collectively veto the implementation of reform in the first period for fear they may lose their jobs. The introduction of elections, however, introduces a starkly different effect at higher average levels of competence. By implementing reforms the cabinet "showcases" its competence and differentiates itself from a (nonactive) opposition. Such behavior whilst undesirable– the expected benefits of reform outweigh the ex-ante risk involved and the degree of learning– is optimal from the cabinet's perspective. Once its politicians are revealed as competent, it will be reelected by a sequentially rational voter.

In a second case, we consider competition between rival executive teams where the shadow cabinet that is is an active participant in the election. We first analyze the simplest case with $q_s = 1$. Then the policy pronouncements of the shadow cabinet are as informative as if they had implemented policy by themselves. Since the details are involved we relegate them to the appendix. In Figure 6 we depict the outcomes. The figures for three cabinet sizes show the parameter space demarcated by three areas: to the left of the dashed line neither team implements the reform policy; to the right of the solid line both teams implement the reform policy; between these lines only the shadow cabinet pronounces in favor of the reform whilst the cabinet implements safe policy.

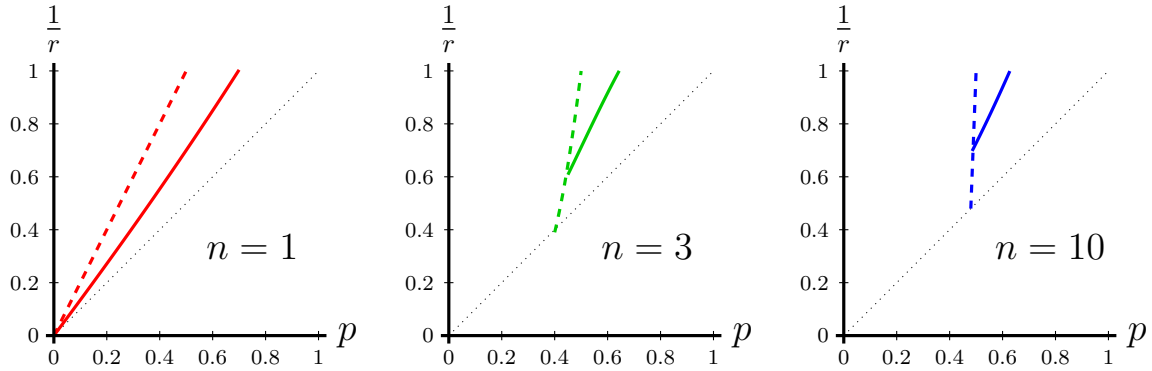

FIGURE 4. Policy Experimentation by a Cabinet that seeks reelection after period 1 with an active Shadow Cabinet. The dotted line indicates $pr = 1$ where expected payoffs from both policies are the same. Each diagram represents different team sizes ($n = 1, 3$, and 10). To the left of the dashed curve, both opposition and incumbent teams adopt the safe policy; in between the dashed and solid curves, the incumbent team adopts the safe policy while the opposition team adopts the risky policy. Finally, to the right of the solid curves and below the dotted line (i.e. when $pr > 1$), both teams adopt the risky policy.

Proposition 5. *Consider a competitive election in which a cabinet with n members, that operates under collective and individual responsibility that is imposed by a party or its leader, seeks reelection against a shadow cabinet that stands on a safe or reform platform. Then relative to the efficient level of risk, implementation of the risky policy in the first period is either too high (at high levels of competence) or too low (at low levels competence).*

When combined with our earlier finding in proposition 3, our result allows us to isolate the mechanism via which learning takes place in cabinet governance. Although that earlier result shows that institutions of cabinet governance induce a hold-up problem that mitigates (indeed eradicates) learning, competitive tension between a cabinet and shadow cabinet restore learning about politicians' competence.Nevertheless, and despite this capacity for learning, it is not straightforward to justify cabinet governance according to such a criterion. As noted in our introductory remarks, a strong defence of cabinet government defined by the imposition of collective responsibility is that "it ensures the will of the majority be put into effect." In considering whether competition between two rival teams can induce the cabinet to induce policies that are optimal from the perspective of a representative voter, our findings provide no support for this line of defence: relative to the optimal level, the level of reforms will be either to low or to high.

7. Optimal Retention in Competitive Elections

In this section we provide a more detailed analysis of the optimal retention rule of the voter. As we have shown, relative to the optimal level, the level of reforms will be either to low (at low levels of competence) or to high (at high levels of competence). These results are premised on voter behavior whereby if the voter is strictly indifferent between electing either team then she selects each with probability one half. Whilst such a tie-break rule is clearly sequentially rational, it is not clear that it is optimal from the voter's perspective. Indeed, precisely because of a voter's indifference, there exists a continuum of equilibria in which the voter breaks the indifference (in different ways). Moreover, it is possible that different equilibria will invoke different levels of implementation of reforms by an incumbent cabinet. And os it follows that, by using a different tie-break rule to the one studied previously, the voter can induce better outcomes. In fact this is so, as we show in this section.

To develop this idea we focus on the simplest possible case, namely setting $q_s = 0$ so that we have a noncompetitive election. First let us consider what it is the voter can hope to achieve. Where the cabinet does not experiment enough with risky reforms, the voter wishes to induce higher levels of

reform implementation. Under these circumstances she could *punish* a cabinet for choosing safe. She can do so by committing to never reelect a cabinet that implements the safe policy. It is sequentially rational for her to do so because when the cabinet plays safe the voter is indifferent between retaining it or electing the shadow cabinet to office.

Where the cabinet over-invests in the risky reforms, the voter wishes to induce lower levels of reform implementation. She can counter the cabinet's incentive to showcase its ability, by *rewarding* safe behavior. This can be achieved by committing to always reelect a cabinet that plays safe. Again this is sequentially rational, since when the incumbent plays safe the voter is indifferent between the incumbent team and a newly elected one. In fact, we show that following this line of reasoning, the voter is able to eradicate inefficiencies in noncompetitive elections.

Proposition 6. *When $pr > 1$ non-competitive elections are strictly welfare improving with respect to the benchmark case. When $pr < 1$ and voters use their votes to both select incumbents and to sanction their performance then noncompetitive elections are equivalent to the benchmark in which the politician is appointed for two terms with no re-election.*

Next we ask whether the voter can design a retention scheme with similar beneficial effects when the shadow cabinet is an active participant.

The left-hand panel in Figure 7 illustrates the welfare comparison between outcomes produced by a cabinet appointed for two terms, which by Lemma 1 are efficient, and those induced by competition between a cabinet and shadow cabinet where $q_s = 1$ where the voter breaks indifference evenly. The light shading indicates the parameter region where competition is welfare enhancing. We note that this includes the entire region where $pr > 1$; this welfare enhancement is due to the fact that in competitive elections the voter is able to vote out of office a cabinet that contains incompetent politicians. It also includes some region of the parameter space where $pr < 1$. The welfare effect is due to learning. When the cabinet chooses the safe policy the shadow cabinet has an incentive to pronounce in favor of the reforms. Doing so allows the voter to learn about the competence of shadow cabinet members and thus make a more informed judgement.

The darker shaded regions, by contrast, are those where competitive elections induce inefficiencies. These occur at both low and high levels of competence for the reasons mentioned earlier in our discussion of proposition 4.

Can these inefficiencies be eradicated when a voter adopts the optimal retention scheme by using the type of strategies discussed above? The right hand-side panel of Figure 7 provides the answer,

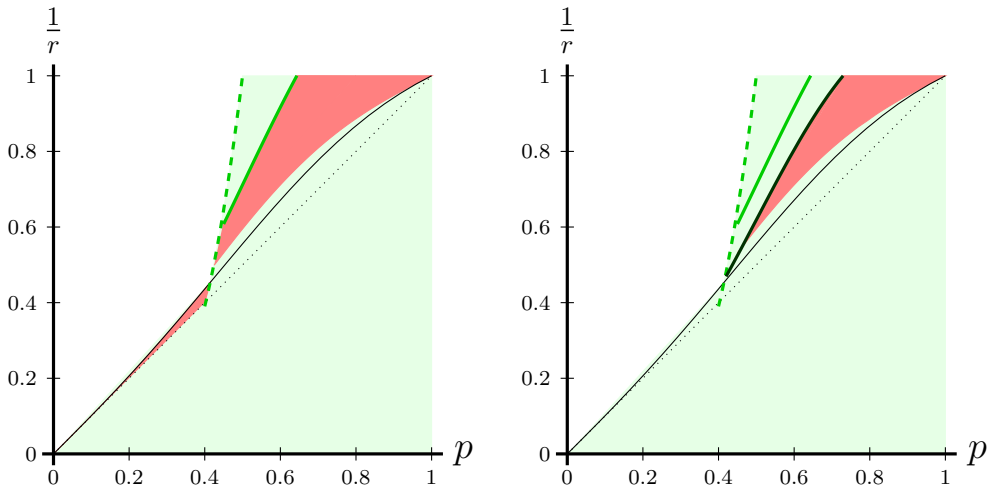

FIGURE 5. Competitive election: Cabinet-vs Shadow Cabinet. The dotted line indicates $pr = 1$ where expected payoffs from both policies are the same. To the right of the solid thin curve, an appointed team adopts a risky platform. To the left of the dashed curve, both opposition and incumbent teams adopt the safe policy; in between the dashed and solid curves, the incumbent team adopts the safe policy while the opposition team adopts the risky policy. Finally, to the right of the solid curves and below the dotted line (i.e. when $pr > 1$), both teams adopt the risky policy. Dark (red) shading indicates inefficiency relative to appointment of the politician. Light (green) shading indicates strictly superior outcomes than appointment. No shading indicates equivalent outcomes as under appointment.

by depicting the same welfare comparison when the voter adopts her best scheme. We observe immediately that, whereas inefficiencies at lower levels of average competence are eradicated, those at higher levels of competence are still present, albeit mitigated to some extent by the voters strategy. In short, even in the optimal scheme, the incentives for the cabinet to showcase by overinvesting in reform remain.

Figure 7 depicts outcomes for a specific cabinet size. We now show how our results depend upon the size of the cabinet.

Proposition 7. *When voters use their votes to both select incumbents and to sanction their performance then the control of incentives eradicates inefficiencies relative to the benchmark appointment case under noncompetitive elections for any team size and under competitive elections for teams composed by single individuals. When teams have strictly more than one member, inefficiencies in competitive elections can not be totally eradicated. In the limit, as the team size grows large, these inefficiencies occur for* $p > \frac{1}{2}$ *and* $pr < 1$.

As we saw earlier a voter can control "showcasing" by rewarding safe behavior by the cabinet in the case where the shadow cabinet is inactive, and so unable to reveal its competence. But a

voter can not eliminate showcasing when there is competition between a cabinet and an active shadow cabinet. Indeed then the inefficiencies that arise due to "showcasing" are increasing in the size of the rival teams. This might seem surprising. The effectiveness of showcasing is limited when the shadow cabinet has a similar strategy that it can adopt. The problem that the voter faces, however, is that it can no longer commit to rewarding the implementation of the safe policy. In particular, when the shadow cabinet pronounces in favor of the reform, and so showcases its talent, then a sequentially rational voter will elect it over a cabinet that plays safe. Alternatively, however, the voter can punish the incumbent cabinet in the event where both the cabinet and the shadow cabinet advocate reform and the voter is indifferent between them. Note however that such indifference implies that both the cabinet and its shadow have the same number of competent members. The likelihood of this event diminishes in the size of the team. The result then follows. A voter can obtain efficient outcomes only in the event where competition is structured between a cabinet and shadow cabinet where the latter is an inactive participant in the electoral contest, or in the where electoral competition is structured between a single incumbent and a challenger. We therefore conclude that whereas electoral competition between a cabinet and shadow cabinet induces learning (by mitigating the free-rider problem) it also introduces inefficiencies.

8. Conclusion

Evidence from several studies suggest that turnover in parliamentary executives is related to a notion of learning about the competencies of politicians. We develop a model of decision-making in teams that allows us to explore whether cabinet governance is in fact conducive to such learning. In our model a politician's type is revealed only when the Cabinet implements reform policies. Critically, however, the Cabinet must collectively agree to the implementation of such reforms. We find that whilst a cabinet can alleviate the free-rider problem inherent in teams, it does so at the expense of a hold-up problem that inhibits learning: members of the Cabinet anticipate that even if successful at implementing reform and thus revealing their competence, further reforms may be vetoed by their colleagues and so are unwilling to make such risky investments. We conclude, therefore, that the central features of cabinet government, namely the consent that is required to the policies implemented by the cabinet, is not conducive to learning.

Nevertheless we show that learning is recovered in a competitive environment where the Cabinet is challenged by a Shadow Cabinet. However, relative to an efficient benchmark, the implementation of reforms is excessive as a cabinet that is (on average) of high competence wishes to showcase its

talents. For such parameter values a representative voter would be able to control an individual decision-maker but is unable to restrain the excessive reforms implemented by a cabinet. Furthermore, the voter would be better off with a team of politicians that is appointed rather than elected.

Our focus on decision-making and learning in a cabinet with collective responsibility thus produces stark findings. We have built our analysis around a situation in which both the cabinet and shadow cabinet choose between the same menu of options where a particular type of policy reveals the competence of the players. We note, however, that our results would continue to hold, albeit in a less clear cut form, so long as there is some action (or set of actions) that a cabinet and its shadow could take that is more revealing of its competence than another. Moreover, and as already noted, our core results would continue to hold under different decision-making rules. The hold up problem that inhibits learning in cabinet governance, for example, is not limited to decision-making under unanimity.

Indeed some of our some of our insights may extend to team decision-making in other environments. In particular our analysis is relevant to situations where agents value their careers as members of a team, where there are some actions that reveal the competence of the team's members, and where the actions taken require the consent of team members. Whilst these features are particularly relevant to political executives that operate under collective responsibility, they exist, albeit in a weaker form, in other organizations. As an example, suppose that a soccer club brings in a new manager in the hope of achieving greater success. This objective is shared by the existing set of players who all want to remain at the club. The manager considers revising the teams tactics, but this involves risk due to uncertainty as to whether his players have the technical expertise and skills necessary to adapt their play. Suppose however that the manager needs the backing of a pivotal section of his dressing-room, perhaps the more senior and established players, to implement these plans. Our analysis suggests that the desire of the players to remain at the club (irrespective of their ability) means that, at best limited, learning can take place. Moreover the introduction of an accountability mechanism– suppose that the manager had the option of firing players revealed as incompetent, or the owner of the club had the option of replacing lock-stock the manager and the players–would lead to suboptimal learning relevant to a benchmark where the team was secure in its position for a fixed period of time.[5]

[5] *Cognoscenti* will recognize the relevance of this parable to a particular English Premiership club, one that happens to be close to one of the authors hearts.

9. Appendix: Formal Proofs

Proof of Proposition 1. In the last period, a reform policy is implemented (successfully) if at least one of its member knows she is competent; in that circumstance the payoff to any agent would be r. It follows that the expected payoff to an agent from everyone implementing a reform platform in period 1 is

$$\sum_{k=0}^{k=n} \binom{n}{k} p^k (1-p)^{n-k} \frac{k}{n} r + \left(1-(1-p)^n\right) r + (1-p)^n$$

where the first term is the average payoff obtained by the team and is equal to pr. The second term is the probability that *at least* one agent is successful in period 1 so the agent obtains a payoff of r in period 2. The last term is the probability *all* agents are not competent thus a safe policy is implemented in period 2 and the payoff is 1. The expected utility of everyone playing safe in period 1 is 2. This implies that in the first best, a team of agents implements reforms only when $\frac{1}{r} < \frac{p+1-(1-p)^n}{2-(1-p)^n}$.

Individual incentives yield free-riding because an agent does not internalize the positive externality that experimentation in period 1 has on the remaining agents. In order to characterize the equilibrium it will be useful to define the expected payoff in period 2 when k citizens experiment in period 1 as

$$f(p,r,k) = \left(1-(1-p)^k\right) r + (1-p)^k = r - (1-p)^k (r-1)$$

The equilibrium has $k \in \{0,1,...,n\}$ citizens experimenting when:

$$\begin{cases} 1 + f(p,r,k) \geqslant pr + f(p,r,k+1) \\ pr + f(p,r,k) \geqslant 1 + f(p,r,k-1) \end{cases}$$

Given that f is an increasing and concave function of k, we have that the in equilibrium there are exactly $\tilde{k}$ agents adopting the reform platform in period 1 where $\tilde{k}$ is the maximum value for which $pr + f(p,r,\tilde{k}) \geqslant 1 + f(p,r,\tilde{k}-1) \Leftrightarrow \frac{1}{r} < \frac{p+1-(1-p)^n}{2-(1-p)^n}$. Note that we should add the further restriction that $\tilde{k}$ should always be smaller or equal than n – obviously there cannot be more agents adopting the reform platform than agents are in the team. □

Proof of Proposition 2. In the last period, a reform policy is implemented (successfully) if all cabinet members know they are competent. It follows that the expected payoff from implementing a reform platform in period 1 is $pr + p^n r + (1-p^n)$; the expected payoff from safe is 2. The unique equilibrium has everyone implementing the reform platform if and only if the former is larger than the latter and so we obtain the stated condition $\frac{1}{r} < \frac{p+p^n}{1+p^n}$. □

Proof of Proposition 3. In the last period, a reform policy is implemented (successfully) if all cabinet members know that they are competent. As opposed to a situation with no replacement we now have that a non-competent cabinet member receives zero utility in period 2. It follows that the expected payoff from implementing a reform platform in period 1 is $pr + p^n r + p\left(1 - p^{n-1}\right)$; the expected payoff from safe is 2. It is immediate to show that no cabinet member will ever want to experiment when $pr < 1$. □

Proof of Lemma 1. With a cabinet in office, we can compute the ideal level of experimentation from the voter's perspective. In this computation we abstract from the individual incentives of politicians to seek reelection but we take into account that the reform policy will be implemented in period 2 only when *all* members of the cabinet know that they are competent. The expected utility of a voter when the cabinet plays safe in period 1 is 2. The expected utility from playing reform is from playing safe

$$\sum_{k=0}^{k=n} \binom{n}{k} p^k (1-p)^{n-k} \frac{k}{n} r + p^n r + (1 - p^n) = pr\left(1 + p^{n-1}\right) + (1 - p^n)$$

By comparison of expected utilities when the cabinet adopts safe or reform in period 1, it follows that the voter wants experimentation in the same circumstances as when the voter appoints a team of officials (with no replacement) into government. That is, when $pr > \frac{1+p^n}{1+p^{n-1}} \Leftrightarrow \frac{1}{r} < \frac{p+p^n}{1+p^n}$. □

Proof of Proposition 4. We find the subgame perfect Nash equilibrium by backward induction. The voter always retains an incumbent cabinet whose average posterior probability of being competent is larger than p. Given the possibility of replacement and that the party (or its leader) will (optimally) replace non-competent members, we have that the cabinet will be reelected if *at least* one member is competent. When the cabinet plays safe, the voter retains or replaces the cabinet with equal probability. When $pr > 1$ it can easily be shown that the cabinet always adopts the reform platform. The interesting case is when $pr < 1$. In that case, a second period cabinet team implements the reform policy if and only if all its member are known to be competent. From the perspective of an individual member of cabinet we know that if she plays reform in the first period then with probability p she obtains r in period 1 and $p^{n-1}r$ in period 2 (probability that all remaining members of cabinet are successful times r). But with probability $1 - p$ she obtains zero in both periods. When a member of cabinet plays safe in period 1 then her expected payoff is $3/2$. It follows that in period 1 the cabinet implements a reform only when $\frac{1}{r} < \frac{p+p^n}{3/2-p+p^n}$. Comparing this last curve to the one in Proposition 2 we can easily show that relative to a situation in which

the cabinet is appointed (and no member can ever be replaced), the adoption of a reform policy in period 1 is less likely for low levels of p and more likely for high levels of p. □

Proof of Proposition 5. We find the subgame perfect Nash equilibrium by backward induction when $pr < 1$ (the opposite case can be trivially solved). The voter re-elects a cabinet when the average posterior probability of its members being competent is above the average posterior of the shadow cabinet. When both teams are equally competent, the voter selects each team with equal probability. Anticipating this voter behavior then the best response of the shadow cabinet to the first period choice of the cabinet is as follows. First, when the incumbent team chooses reform and at least one cabinet member is successful, the shadow cabinet has a positive probability of being elected and so receives a positive payoff only when she also chooses reform. Second, when the cabinet plays reform and all its members are unsuccessful or when it plays safe the shadow cabinet knows that it will be elected with probability $1/2$ if playing safe; instead, the expected utility of any shadow cabinet member when playing reform is $p\left(p^{n-1}r + \left(1 - p^{n-1}\right)\right)$ (note that we are taking into account that a shadow cabinet member that shows to be non-competent is replaced). It follows that the shadow cabinet adopts a reform platform when the cabinet's average posterior is p and $\frac{1}{r} < \frac{p^n}{0.5-p+p^n}$. Finally we solve for the cabinet's optimal first period choice when anticipating the opposition and voter's best responses. There are two cases to solve for. When $\frac{1}{r} > \frac{p^n}{0.5-p+p^n}$, the opposition best response to a cabinet that plays safe (or one that plays reform but everyone is non-competent and is replaced) and a cabinet member expected utility when playing safe is $1 + \frac{1}{2}$, and when implementing the reform policy is

$$\begin{aligned} &pr + p \cdot \Pr\left(\text{being reelected} \mid \text{i is competent}\right) + p \\ &\quad \cdot \Pr\left(\text{all cabinet members are successful} \mid \text{i is competent}\right) \cdot \left(1 - \frac{p^n}{2}\right) \cdot (r-1) \end{aligned} \tag{1}$$

The expression for the second term $\Pr\left(\text{being reelected} \mid \text{i is competent}\right)$ is not simple so we consider upper and lower bounds for this expression. Doing so we show that the expression we are interested in is bounded below and above by $\frac{1}{2}$ and $\left(1 - \frac{p}{2}\right)$. Note that we are trying to characterize

$$\Pr\left(X = k \mid X \geq 1\right) \cdot \left(\Pr\left(Y < k\right) + \frac{1}{2}\Pr\left(Y = k\right)\right)$$

where X and Y are two independent random variables following a binomial distribution with parameters n and p, respectively. This expression is decreasing in n given that as n increases,

knowing that X is at least 1 tends to have less weight; in the limit as n grows to infinity, the expression is equal to $\frac{1}{2}$.

Using these bounds it can now easily be shown that the expression in (1) will always be smaller than $1+\frac{1}{2}$. In other words, the cabinet always implements the safe policy when the shadow cabinet's best response to safe is safe.

When $\frac{1}{r} < \frac{p^n}{1/2-p+p^n}$, the shadow cabinet always stands on a reform platform regardless of the cabinet's action. A cabinet member's expected utility when playing safe is $\left(1+\frac{1}{2}(1-p)^n\right)$ and when implementing a reform platform is equal to the expression (1) above. We conclude that the cabinet will adopt a reform policy whenever $\frac{1}{r} < \frac{p^n}{1/2-p+p^n}$ and $\frac{1}{r} < f(p)$ where $f(p)$ belongs to the following interval:

$$\left[\frac{2p+p^n(2-p^n)}{2+(1-p)^n-p+p^n(2-p^n)}+\frac{2p+p^n(2-p^n)}{2+(1-p)^n-p(2-p)+p^n(2-p^n)}\right] \tag{2}$$

It can be easily shown that for $pr<1$, $f(p)$ increases with n. In other words, as n grows large both the cabinet and its shadow adopt a safe policy for $p<\frac{1}{2}$ and they both implement/adopt a reform policy for $p>\frac{1}{2}$. □

Proof of Proposition 6. We claim that the voter can eradicate inefficiencies under non competitive elections ($q_s=0$). Recall that under a noncompetitive election the voter can never commit to anything other than electing the team whose average posterior probability of being competent is highest. This implies that she can only provide *incentives* when she is strictly indifferent between retention or not of the cabinet. From proposition 4 we know that inefficiencies arise in the parameter region where $pr<1$ and when

(i) p is small: for some parameter values, the voter would like the cabinet to adopt the reform policy when instead it takes the contrary action. In this circumstance the voter could *punish* safe behavior by never reelecting a cabinet that plays safe. By doing so a cabinet member obtains a payoff of 1 when playing safe but a payoff equal to the expression in (1) when choosing reform. It is immediate to show that a cabinet member will always prefer to adopt the reform policy in the area in which the voter would prefer her to do so.

(ii) p is large: for some parameter values, the voter would like the cabinet to not adopt the reform policy. Under these circumstances the voter could *reward* safe behavior by reelecting with certainty a cabinet that plays safe. Once again, straightforward calculations reveal that such incentives are enough for the voter to recover her preferred level of experimentation. □

Proof of Proposition 7. From proposition 5 we know that there are two areas where inefficiencies arise. When p is small, the voter would like the cabinet to adopt the reform policy when it is not doing so. In this case, the same incentive scheme described in the proof of Proposition 4 could be used to recover efficient levels of experimentation (*punish* safe behavior). When p is large, the cabinet takes too many risks. The voter can implement the right incentives when indifferent: she can reward safe behavior by reelecting the cabinet when playing safe and when the shadow cabinet replaces all of its members (recall that in the area we are concerned about, the shadow cabinet always adopts a reform platform); and she can punish risk taking behavior when the cabinet adopts a reform policy and its posterior average probability of being competent coincides with the one of the shadow cabinet. These are the only sequentially rational punishments the voter can adopt. Note that the likelihood of the voter being indifferent and able to appropriately reward and punish the cabinet decreases with n. Given these incentives we can compute the parameter values for which the cabinet will adopt a reform policy when comparing the expected payoff of both policies. For notational simplicity we write these expressions for the case in which $n = 2$, although the analysis extends easily,

$$1 + (1-p)^2 > pr + p^2(1-p^2)r + p(1-p)(1-p)^2$$

where the left-hand-side of the inequality denotes the expected utility from adopting a safe policy and the right-hand-side the one from adopting a reform policy.

We have that the cabinet adopts a reform policy when the voter uses an optimal retention rule only if $\frac{1}{r} < \frac{p+p^2(1-p^2)}{1+(1-p)^2-p(1-p)(1-p)^2}$. Previously (in the absence of optimal incentives) we showed that the cabinet adopted a reform policy when $\frac{1}{r} < \frac{p+p^2(1-\frac{p^2}{2})}{1+\frac{1}{2}(1-p)^2-p(1-p)((1-p)^2+p(1-p))}$. A comparison of these expressions yields the result that the voter is able to constrain risk taking by the incumbent, but whenever $n > 1$, she can no longer recover the level of experimentation of the (constrained) first best. As n grows large, both the right-hand-side of the restrictions above tend to the same value and so as n grows large the voter can do no better than break indifference evenly. □

References

Aghion, P., P. Bolton, C. Harris, and B. Julien (1991): "Optimal Learning by Experimentation," *Review of Economic Studies*, 58(4), 621–54.

Alesssina, A., and G. Tabellini (2007a): "Bureaucrats or Politicians? Part I: A single Policy Task," *American Economic Review*, 97, 169–179.

——— (2007b): "Bureaucrats or Politicians? Part I: A single Policy Task," *Journal of Public Economics*, 92, 426–447.

Ashworth, S., and K. W. Shotts (2010): "Does informative media commentary reduce politicians' incentives to pander?," *Journal of Public Economics*, 94(11-12), 838–847.

——— (2011): "Challengers, Democratic Contestation, and Electoral Accountability," working paper.

Banks, J. S., and R. K. Sundaram (1990): "Incumbents, Challengers, and Bandits: Bayesian Learning in a Dynamic Choice Model," Rochester Center for Economic Research, Working Paper.

——— (1993): "Adverse Selection and Moral Hazard in a Repeated Elections Model," in *Political Economy: Institutions, Competition, and Representation*, ed. by W. A. Barnett, M. J. Hinich, and N. J. Schofield, chap. 12, pp. 295–312. Cambridge University Press.

——— (1998): "Optimal Retention in Agency Problems," *Journal of Economic Theory*, 82(2), 293–323.

Barro, R. (1973): "The Control of Politicians: an Economic Model," *Public Choice*, 14(1), 19–42.

Berlinski, S., T. Dewan, and K. Dowding (2010): "The Length of Ministerial Tenure in the UK, 1945–97," *Journal of Politics*, 74(2).

——— (2012): *Accounting for Ministers: Scandal and Survival in the British Cabinet*. Cambridge University Press, Cambridge.

Besley, T., and A. Case (1995): "Does Electoral Accountability Affect Economic Policy Choices? Evidence from Gubernatorial Term Limits," *Quarterly Journal of Economics*, 110(3), 769–798.

Bolton, P., and C. Harris (1999): "Strategic Experimentation," *Econometrica*, 67(2), 349–374.

Callander, S. (2011a): "Searching and Learning by Trial and Error," *American Economic Review*, 101(3), 2277–2308.

——— (2011b): "Searching for Good Policies," *American Political Science Review*, 105(4), 643–663.

Callander, S., and B. Hardstad (2012): "Policy Experimentation in Federal Systems," working paper.

Cox, G. (2011): "War, Moral Hazard, and Ministerial Responsibility: England After the Glorious Revolution," *The Journal of Economic History*, 71(1), 133–61.

Dewan, T., A. Galeotti, C. Ghiglino, and F. Squintani (2011): "Information Aggregation and Optimal Executive Structure," Working paper.

Dewan, T., and R. Hortala-Vallve (2012): "Policy Learning and Elections," working paper.

Dewan, T., and R. Hortalla-Valve (2011): "The 3 **A**s of Government Formation: Appointment, Allocation, and Assignment," *American Journal of Political Science*, 55(3).

Dewan, T., and D. P. Myatt (2007): "Scandal, Protection, and Recovery in the Cabinet," *American Political Science Review*, 101(1), 63–77.

——— (2010): "The Declining Talent Pool of Government," *American Journal of Political Science*, 54(2), 267–286.

Ferejohn, J. (1986): "Incumbent Performance and Electoral Control," *Public Choice*, 50(1), 5–25.

Fu, Q., and M. Li (2010): "Policymaking with Peputation Concerns," Working Paper.

Hirsch, A. V. (2011): "Experimentation and Persuasion in Political Organizations," Paper presented to the Annual Meetings of the American Political Science Association, August 2011.

Holstrom, B. (1982): "Moral Hazard in Teams," *Bell Journal of Economics and Management Science*, 13, 324–40.

Huber, J. D., and C. Martinez-Gallardo (2003): "Cabinet Turnover in Parliamentary Democracies," paper presented to the annual meetings of the American Political Science Association.

——— (2008): "Replacing Cabinet Ministers: Patterns of Ministerial Stability in Parliamentary Democracies," *American Political Science Review*, 103(2), 169–180.

Indridason, I., and C. Kam (2008): "Cabinet Reshuffles and Ministerial Drift," *British Journal of Political Science*, 38(4), 621–656.

Lizzeri, A., and N. Persico (2009): "Electoral Incentives, Political Risk-Taking and Policy Reform," in *The Political Economy of Democracy*, ed. by E. Aragones, C. Bevia, L. H., and N. Schofield. Fundación BBVA.

Majumdar, S., and S. W. Mukand (2004): "Policy Gambles," *American Economic Review*, 94(4), 1207–1222.

Marshall, G. (1984): *Constitutional Conventions: The Rules and Forms of Political Accountability*. Oxford University Press, Oxford.

Strøm, K. (1985): "Party Goals and Government Performance in Parliamentary Democracies," *American Political Science Review*, 79(3), 738–754.

Strulovici, B. (2010): "Learning While Voting: Determinants of Collective Experimentation," *Econometrica*, 78.

Strumpf, K. S. (2002): "Does Government Decentralization Increase Policy Innovation," *Journal of Public Economics*, 4, 207–41.

Turner, D. (1969): *The Shadow Cabinet in British Politics*. Routledge and K. Paul.

Volden, C., M. M. Ting, and D. P. Carpenter (2008): "A Formal Model of Learning and Policy Diffusion," *American Political Science Review*, 102(3), 319–332.

Williamson, O. E. (1992): "Transaction-costs Economics: The Governance of Contractual Relations," *Journal of Law Economics and Organizations*, 22(2), 233–62.