Political Science and Political Economy Working Paper

Department of Government

London School of Economics

No. 1/2012

Policy Learning and Elections

Torun Dewan (LSE)

Rafael Hortala-Vallve (LSE)

THE LONDON SCHOOL OF ECONOMICS AND POLITICAL SCIENCE

Policy Learning and Elections*

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 compare performance of politicians under appointment and election in a situation where learning by trial and error can yield better policy outcomes. We model an incumbent who chooses between a "safe" option with a known payoff and a "risky" policy that yields higher or lower payoffs depending on the policymaker's type. We consider two types of elections: in noncompetitive elections voters learn about the incumbent's type and compare expected payoffs from retention against dismissal and replacement by another representative agent; in competitive elections, an opponent campaigns on a "safe" or "risky" platform, voters evaluate the type of both incumbent and challenger, and choose between the two alternatives. Politicians prefer holding office but otherwise politicians and voters are symmetric with respect to policy preferences and information. An appointed agent will engage in "risky" policies even when expected first period returns are relatively low. Relative to this efficient benchmark, elected politicians either under invest in the risky option due to a "fear of failure" or over invest by "gambling on success". Despite this, competitive elections can induce efficiency gains. Moreover, when voters use their votes to simultaneously select high performers and provide incentives, competitive elections are unambiguously superior to appointment. We test our hypothesis in a laboratory experiment that yields support for our findings.

*For helpful comments we thank Georgy Egorov, Fabrizio Facundo, Francesco Giovonnoni, Stuart Jordan, Ben Lockwood, Leandro de Magalhaes, Ken Shepsle, Jim Snyder, Francesco Squintani, Alexey Zaharov, participants at the Workshop on "Institutions, Individual Behavior and Economic Outcomes" at the University of Sassari, at the Workshop of Sonderforschungsbereich 884 "Political Economy of Reforms" at the University of Mannheim, and audience members at talks given at the Higher School of Economics, Moscow State University, the Department of Economics at Bristol University, the Department of Economics at Warwick University, and at the Annual Meetings of the American Political Science Association, Seattle, 2011.

1. Introduction

Should decision-makers be elected or appointed? The question is pertinent to recent cases in Greece and Italy where implementation of reforms have been entrusted to appointed rather than elected officials. Indeed the current Italian government formed on November 15th, 2011, does not contain a single elected politician. The standard model of agency shows that elections allow for the screening of politicians and the provision of incentives for better performance. A newer literature (Canes-Wrone, Herron, and Shotts, 2001; Maskin and Tirole, 2004), however, suggests instead that elections "may induce those officials to pander to public opinion and put too little weight on welfare" (Maskin and Tirole (2004),p1034). Such pandering occurs (in these models) under assumptions of asymmetric preferences and information between voters and office-seeking politicians. We explore a different policy environment: both politicians and voters have the same policy preferences but, *a priori*, are uninformed as to which actions are in their best interests; they learn from policy outcomes related to the actions taken by politicians.

Such learning on the job is an important aspect of policymaking in business and in politics: the CEO of a company evaluates current decisions on the basis of past outcomes; a politician with executive power learns about the policy environment and her own abilities via her track record of policy implementation. Perhaps the most well known era of such policy learning is the series of policy experiments that are collectively known as the New Deal, including such landmark measures as the National Industrial Recovery Act, the Agricultural Adjustment Act, and the establishment of the Tennessee Valley Authority.[1] Many contemporary accounts of this period note that these policies were at best weakly founded on economic principles and were instead implemented on a trial and error basis. Indeed Roosevelt himself said of the times that "the country needs and, unless I mistake its temper, the country demands bold persistent experimentation. It is common sense to take a method and try it, if it fails admit it frankly and try another."

Learning depends upon the willingness of a policymaker to experiment rather than stick with tried and tested options. Tim Harford, author of *Adapt: Why Success Always Starts with Failure*, highlights differences in risk-taking across sectors. In the business world, he argues, success is built on previous failure. By contrast, he asks, "in politics where are the bad ideas that have been tested,

[1]Beyond US national politics there are many examples.Hall (1993) amongst others has looked at the shift from Keynsian to Monetarist economics in the UK during the 1970s and 80s as an example of policy experimentation and learning; in China the return to household from collective farming first initiated in Anhui Province and then elsewhere has been studied as policy learning by Cao, Qian, and Weingast (1999). At the local level in the United States several localized experiments have been implemented and extensively researched, amongst them Nevada's decision to legalize prostitution and the health care plan in the state of Oregon.

found wanting, and replaced with something better?"[2] From a political economy perspective the question of interest is whether elected politicians can deliver efficient levels of risk-taking. Whilst the implementation of policy on a trial and error basis might appear a pragmatic course of action [3] experiments may not deliver welfare improvements and instead produce costs that are born by the public.[4] Then ex-post we would prefer the decision maker had stuck with original policy. Ideally experimentation should be neither too high or to low, but efficient from the public view.

Can elected politicians produce efficient levels of risk or should such decisions best be entrusted to appointed officials? We develop a simple model to explore this question. A politician is faced with implementing one of two policies: the first "safe" option has a known payoff (to voters and the politician alike); whilst the other "risky" option can in some situations yield a higher payoff relative to the safe option, though may be worse ex-post. Should the politician implement the "risky" option then its success depends upon the incumbent's type: the politician may be competent or not, and at the time of making the initial decision is unaware of her abilities. After the policy is implemented voters (and the incumbent) learn whether the policy has been a success and therefore also about the incumbent's type. In any equilibrium of our model, the policies implemented by the incumbent and voters' final decisions are sequentially rational.

We first analyze a benchmark where the policy decision is taken by a representative agent who does not face re-election. When serving two terms rather than one, an agent may take a risky policy in the first period though her expected first period payoff is *lower* than obtained when following the safe option. In doing so she learns her type and so she is better informed when taking the second period action.

Then we introduce the standard model of electoral competition: the incumbent can be removed after the first period and replaced with a politician of an unknown type. This leads to experimentation by an incumbent that is either (weakly) higher or lower than efficient levels. When the prior probability that she is competent is low, a politician does not want to reveal her incompetence by taking a risky option; fearing failure she instead plays safe. By contrast, when the prior probability she is competent is high then a politician may be *too* willing to take the risky option: anticipating success she experiments when a representative agent would not; the electorate will always reelect an incumbent on a proven track record and so she "gambles on success." These findings reflect

[2]http://www.freakonomics.com/2011/05/10/why-is-failure-a-sign-of-a-healthy-economy
[3]as per the maxim "there is no such thing as a failed experiment- only unexpected outcomes"
[4]For example DeCicca and Smith (2011) assess the negative consequences of the introduction of "dual entry"-the possibility of preschool entry at different dates in British Columbia.

the stark nature of the retention choice: a rational voter cannot commit to reelecting a politician whose policy is not successful; nor to removing one who has been successful.

Next we ask how "gambling on success" and "fear of failure" are related to competitive aspects absent in the standard model. In a 'competitive' election, after the first period has elapsed, an opponent chooses which policy– "safe" or "risky"– to campaign on. If she chooses the latter, then voters (and the opponent) learn with some probability whether she is competent or not. If she chooses safe then there is no learning. The introduction of an opponent opens up interesting strategic possibilities: the opponent can benefit by showing "wisdom after the act" when the incumbent's policy experiment fails (by advocating safe options); or, if the incumbent plays safe, she might gain by advocating risky policies that suggest boldness relative to the incumbent. In this institutional environment the retention of the policy maker is tied to her performance and to the perceived competence of her opponent. However, we find that, as before, first period investment in the risky option by the incumbent is either higher or lower than in the benchmark. Moreover, it is always lower than is the case with noncompetitive election.

These results might suggest that elections yield inefficient outcomes. This is in fact not the case. Indeed competitive elections lead to efficiency gains at lower levels of competence. The reason is that voters learn from the opponent's policy pronouncements and thereby attain better outcomes. In particular, when an incumbent chooses the "safe" policy in the first period and her opponent's best response is to advocate that risky policies be implemented then this always benefits voters: in making the case for the risky policy, voters learn about the opponent's type (and thereby anticipated policy outcomes) *without* incurring the fixed costs of experimentation by the incumbent.

Nevertheless, inefficiencies remain under competitive elections when incumbents gamble on success. Such inefficiencies could perhaps be mitigated when voters punish such behavior in equilibrium. We show that such strategic punishment is sequentially rational and eradicates inefficient behavior. Thus in the policy environment we analyze, contrary to conventional wisdom (Fearon, 1999), voters can use their votes to simultaneously select politicians and provide incentives. And, once we consider the role that selection and incentives play then elections can be welfare enhancing. These welfare effects are apparent, however, only when considering competitive interaction between an incumbent and a challenger. Thus the inefficiencies of elections highlighted in the models cited earlier may be due to their stark focus on the retention choice. Indeed this screening effect– whereby incompetent or bad politicians are weeded out– does not distinguish elections from other accountability mechanisms common to other institutions. For example, chief executives are held

accountable by management boards for their company's risk management and, in non market economies, party officials can be held responsible for policy failures and dismissed by their party superiors. It is the competitive aspect that distinguishes elections as an accountability mechanism – a politician, unlike the business executive or party official, faces periodic evaluation by an electorate who assess her record of achievement *against* the pronouncements of an opponent– and that can lead to welfare improvements due to learning.

The predictions of our model run counter to the so-called "certainty effect" in behavioral economics. In their seminal article Kahneman and Tversky (1979) provide numerous examples where subjects underweigh the expected gains from a risky outcome relative to a sure thing when the prospect of winning the prize is large; and underweigh the expected loss of a risky option relative to a sure thing when the prospect of winning the prize is small. In their words (Kahneman and Tversky (1979), p269) "in the positive domain the certainty effect contributes to a risk averse preference for a sure gain over a larger gain that is merely probable. In the negative domain, the same effect leads to a risk seeking preference for a loss that is merely probable over a smaller loss that is certain." By contrast, in our model the incumbent overweighs the positive gain from the risky policy relative to the safe option when she has a high chance of securing the prize, whereas she overweighs the loss from the risky policy relative to the safe option when she has a low chance of winning. Our hypothesised effects can then be identified in a suitably controlled environment.

We explore the strategic aspects of our model in a laboratory experiment that reproduces the incumbent's first period choice and the effects of noncompetitive and competitive elections. Our findings are supportive of our main claim. In particular, when subjects' anticipated payoffs reflected optimal behavior of voters and a strategically rational opponent then their first period behavior accorded with subgame perfect Nash strategies: they under or over invested in the risky option in accordance with the predictions of our model.

In section 2 we introduce our model. In section 3 we provide our benchmark analysis of policy learning with an appointed decision-maker. In section 4 we analyze non-competitive and section 5 competitive elections. Section 6 provides a comparison of efficiency with respect to the benchmark, and between elections. In section 7 we analyze incentives and selection and conclude our comparative evaluation of efficiency and we discuss our insights with respect to incumbency advantage. In section 8 we discuss our experiment. We conclude by summarizing our findings and exploring future avenues of research. We begin by first discussing our contribution to the relevant literature.

2. Related Literature

We contribute to a small (but growing) formal literature on policy learning: Volden, Ting, and Carpenter (2008) analyze a situation where governments learn from their own experiences and those of other governments; Callander (2011a) analyzes learning by trial and error in a business environment and Callander (2011b) learning in a spatial model of elections. In an earlier paper Callander (2008) analyzes learning in a bureaucratic-sponsor relationship.

We explore learning within the context of the two-armed bandit model, used in policy analysis by Aghion, Bolton, Harris, and Julien (1991), adapted by Strumpf (2002) to look at policy innovation and its relation to government decentralization, and by Strulovici (2010) to analyze experimentation by groups of decision-makers. Our focus on electoral competition within this context relates our paper to Banks and Sundaram (1990); they analyze 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. In 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. The policy environment in our model is closest to that used by Lizzeri and Persico (2009) who study the impact of different electoral systems on risk control.

Our model contributes to a vast literature on electoral accountability that goes back to the models of Barro (1973) and Ferejohn (1986), developed further by Banks and Sundaram (1993) and Banks and Sundaram (1998). The key idea in these classic models is that such accountability leads to better outcomes. A recent strand of this agency literature has, however, reached starkly different conclusions: Canes-Wrone, Herron, and Shotts (2001) and Maskin and Tirole (2004) model a situation of asymmetric information where politicians are perfectly informed as to the policy choice that is in voters' best interests but may not share voters' preferences. Inefficiencies owing to electoral concerns arise due to what the authors term "pandering": politicians implement populist measures in order to appear aligned with voter preference.[5] Similar inefficiencies arise in our policy environment with symmetric information and shared preferences–specifically, and using the parlance of Maskin and Tirole, politicians and voters are congruent and are symmetric in their ignorance prior to policies being chosen–and so speak to the robustness of those earlier insights. However, we show that these results are overturned in a competitive environment with an active challenger: then elections are superior to appointment.

[5]Extensions of this framework have looked at different aspects of elections and constitutional design (Fox, 2007; Fox and Stephenson, 2011).

The analysis of competitive elections between an incumbent and an opponent relates our work to Ashworth and Shotts (2011, 2010). In Ashworth and Shotts (2010) the media introduces a commentary on the incumbent's policy record before an election is held. The authors show that whilst this can sometimes attenuate pandering, it can also exaggerate such behavior. In Ashworth and Shotts (2011) an incumbent engages in costly information acquisition before choosing a policy. The role of the challenger is to assess and criticize the incumbent's platform. When such claims are verifiable, voters can use retrospective voting strategies to sanction incumbents. This induces incentives for better performance. By contrast our focus is on selection rather than sanctioning. However in an extension to our basic framework we show that voters can do both and so eradicate inefficiencies associated with elections in our policy environment. The fact that voters can simultaneously use incentives and selection runs contrary to claims made by Fearon (1999). An extensive treatment of this issue is provided by Ashworth, Bueno de Mesquita, and Friedenberg (2011).

Our work offers new insights to incumbency advantage, studied by Ansolabehere and Snyder (2002) and Zaller (1998) amongst others. In existing studies this advantage refers to any increase in vote share due to the status of incumbency– perhaps due to campaign spending (see Ashworth (2006) and Erikson and Palfrey (2006)), or uncertainty over candidate quality (Samuelson, 1987; Gordon, Huber, and Landa, 2007). Even tight causal analysis of the (partisan) incumbency advantage– as provided by Lee (2008)–can not distinguish between these and other channels. We operate a strict definition of incumbency advantage as a systematic pro-incumbent bias that exists even once all uncertainty over type is resolved and show that voters (optimally) exercise bias when ex-post indifferent between candidates. The latter bias occurs when incumbents are on average of high competence. Relatedly, Gordon and Landa (2009) show that common sources of incumbency advantage, such as perceived competence, may in fact work in a challenger's favor.

Finally, our analysis shows that longer term length induces efficient learning. This relates to the literature on term limits in the US states: Carey, Niemi, Powell, and Moncrief (2006) use survey data to explore the effect of term limit imposition on selection and legislative behavior; Besley and Case (1995) highlighted the relatively poor fiscal performance of term-limited US governors between 1950 and 1987, with this effect shown to be robust to a wider span of years in the data Johnson and Crain (2004) and controls for other temporal effects such as experience (Alt, Bueno de Mesquita, and Rose, 2011). Relatedly, Ferraz and Finan (2009) show evidence of less corruption in Brazilian municipalities where the mayor can be re-elected; and Dal Bó and Rossi (2008) show evidence that longer time horizons leads to better performance of Argentinian legislators.

3. A Model of Learning and Elections

We model the strategic interaction between politicians $j \in \{i, o\}$, where i is the incumbent and o her opponent, and a representative voter. Each politician is either competent or not and the prior probability that any politician is competent is $p_j \in (0, 1)$. We assume that the politician's type does not vary with time and that, prior to the taking of any actions, politicians do not know their type or that of the other politician.

The policy environment is such that in each of two periods, one of two policies can be adopted: the first is a "safe" policy that always yields a payoff of 1 to the politician who implements it and to voters regardless of the politician's type; the second is a risky policy. The outcome from implementation of the risky policy can either be a success or a failure. If the politician who implements the risky policy is competent then a success is obtained with probability q_i where $q_i \in [0, 1]$. This yields a payoff of $r > 1$ to the politician and to voters. If the politician who implements the risky policy is not competent then the outcome is failure. This yields a payoff of 0 to both politicians and the voter.

The incumbent chooses the first period policy. After the policy outcome is observed an opponent chooses a policy either "safe" or risky". Although the action set is the same for both politicians there is a key difference between the opponent and the incumbent: whilst the latter stands on an observed record of implemented policy, and will be judged by voters accordingly, the former does not. Instead the opponent can only reveal her competence via her campaign pronouncements. We model this by assuming that should the opponent choose the risky policy she reveals her competence with probability $q_o \in [0, q_i]$. Voters then evaluate the type of both incumbent and challenger before electing one or other in a winner-take-all contest.

Although our model specification is sparse we draw readers attention to several of its key features.

First the payoffs describe a scenario where politicians are both policy-seeking and office-seeking. These twin concerns are captured in a parsimonious way by assuming that politicians receive a policy payoff only when in office. During such times their payoffs are perfectly aligned with those of voters. However, when out of office a politician's payoff is 0. This induces a misalignment in the payoffs of politicians and voters since the politician would like to retain (or attain) office and this feature provides the strategic tension between our players. Note also that the incumbent (and voter) bears the cost when the outcome of the risky policy is failure – they forego the safe payoff–whilst the opponent always obtains 0 irrespective of the incumbent's action.

Second, policy outcomes involve an interaction between the policy implemented and the politicians type, so voters and politicians learn from past experience. Specifically, voters observe policy outcomes after the first period policy has been implemented and so learn about the incumbent's type when she takes the risky option, but learn nothing when she plays safe. The stark specification of a politicians payoffs and the correlation between type and outcomes insures that we obtain crisp results but is without loss of generality.

Third the difference between q_i and q_o captures an important element in the competition between an incumbent and her opponent. The incumbent stands on an observed policy record: in the event where she takes the risky policy then her record is one of success or failure. By contrast her opponent stands on an untested policy platform: she is judged by her words rather than her deeds. With this in mind, and with no corresponding loss of generality, we can set $q_i = 1$ so that there is perfect correlation between the incumbent's type and the policy outcome if she chooses the risky policy. Doing so we can then build our analysis on several interesting cases.

In the first we set $q_o = 0$. Then, whereas voters evaluate the incumbent on her record, any pronouncement by the opponent is dismissed by the electorate (i.e. opponent's pronouncements are uninformative). In this institutional environment the retention of the incumbent is tied solely to her performance; the challenger is irrelevant. This yields a standard model of elections in which they serve as a mechanism for selecting competent politicians. As already noted, this selection mechanism operates in politics as in other walks of life. For example, a company executive can also be fired by the management board when risky ventures do not yield success. Selection provides a mechanism by which the policymaker is held directly accountable for risk management. We refer to this model –in which the opponent plays no role and the game is played between the incumbent and a voter who observes her performance – as one of a "noncompetitive election".

In all other cases we set $0 < q_o \leq q_i = 1$. Then there is a role for the opponent since voters can infer her competence via her campaign. However, in setting $q_o \leq q_i = 1$, the burden of proof is generally harder for the opponent for she must rely on campaign statements rather than on an established record of achievement (or failure). In the (degenerate) case of $q_o = q_i = 1$ the words of the challenger speak as loudly as the deeds of the incumbent. Thus a challenger can, during the campaign, make the case for the risky policy that is as convincing *as if* she had herself implemented the policy. We refer to this as a perfectly competitive election and we make use of this case to show that some of our results, obtained for a noncompetitive election continue to hold even in the most

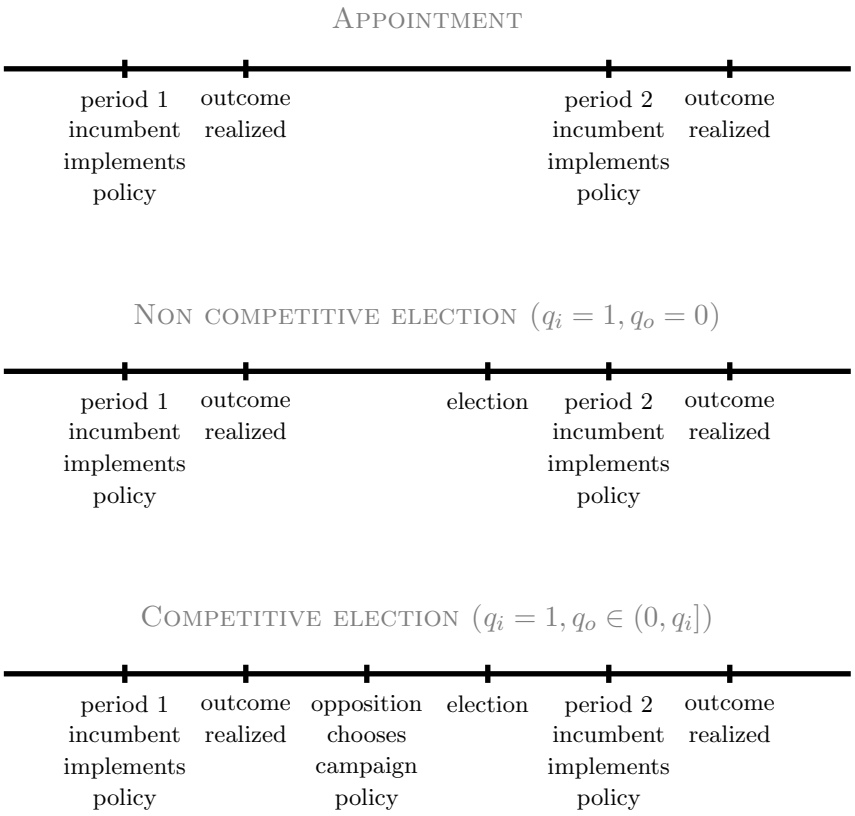

Figure 1. Timeline for Incumbent Evaluation

competitive environment we can consider. We compare elections, whether competitive or not, with a benchmark model in which the incumbent is appointed for two periods.

Figure 3 describes the corresponding time-line for the different cases discussed above. In our benchmark case of appointment the incumbent chooses a policy in both periods with no intervening election. Under a noncompetitive election ($q_o = 0$) then by construction the opponent is irrelevant and we focus solely on the incumbent's action: she chooses a policy and then voters observe the policy outcome before choosing whether to retain or replace her with an opponent of an unknown type. Whomsoever the voter chooses then implements second period policy and the game ends. In a competitive election ($q_o > 0$), by contrast, after the incumbent's policy is chosen and outcomes are revealed she is challenged by an opponent. Voters then evaluate the type of both incumbent and challenger before electing one or other in a winner-take-all contest. The winner of that contest implements the second period policy and the game ends.

We analyze each case in turn, but first discuss voter behavior across cases. Voters condition their behavior on different events that inform their beliefs over politicians competence. Given

the symmetry in preferences their optimal strategy is in selecting the politician who given their posterior beliefs is most likely to be competent. If indifferent, we assume the voter re-elects the incumbent with probability 1/2; later in section 8 we revisit the simplifying assumption.

In a noncompetitive election the incumbent reveals herself to be competent by taking the risky action and succeeding, and incompetent when failing. Likewise in a competitive election. The difference in the latter is an opponent who can also distinguish her type by taking the risky option. We describe standard voter behavior in these settings. In noncompetitive elections the voter always re-elects a competent politician, and selects the opponent if the incumbent reveals herself as incompetent; otherwise –in the event where the incumbent plays safe and so does not reveal anything about her type– the incumbent is re-elected with probability 1/2.[6] In a competitive election if the incumbent reveals herself to be competent then she is elected when her opponent reveals herself to be incompetent or plays safe; if, however, the opponent also reveals herself as competent then the incumbent is elected with probability 1/2. If the incumbent chooses the safe policy then she is elected with probability 1 if her opponent reveals herself to be incompetent; with probability 1/2 if the opponent plays safe; and zero probability if the opponent reveals herself as competent.

4. Policy Learning Under Appointment

To establish a benchmark we first evaluate policy learning in a world without elections. Recall that, given our setup, the strategic tension between the politician and the voter is due to the desire of the latter to attain office and to remain there. Otherwise voters and politicians' incentives are perfectly aligned. It follows that the analysis of our model without elections establishes the efficient benchmark with respect to risk taking: then the official policymaker is simply a representative agent.

If the agent serves in office for a single term only then it follows immediately that she will adopt the risky decision only when $pr > 1$. We then consider efficient risk control in a two period model without reelection after the first term has elapsed. In the first period the agent decides whether to adopt a safe or risky policy. After observing the period 1 outcome which, in the event where she took the risky option, reveals her type, she then decides once again whether to adopt the safe or risky option. The payoff of this two period model is the sum of each period's model (we assume there is no discount factor). A straightforward calculation provides our first result:

[6]The election of the alternative when the incumbent chooses the risky policy and fails is perhaps not immediate. Since a politician who fails will always choose the safe option in the second period the voter is strictly indifferent. We assume that indifference is broken in the opponent's favor: this is sequentially rational if the game continues to a third period and so holds if voters believe there is some (however small) probability that the game continues.

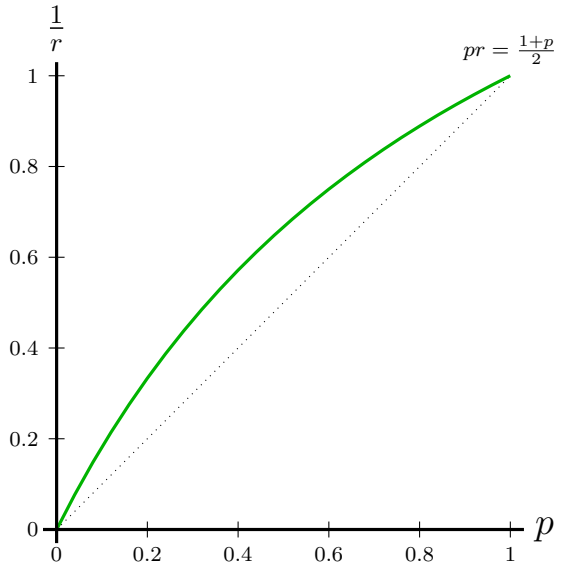

FIGURE 2. Policy Learning by Appointment. The dotted line indicates $pr = 1$ where expected payoffs from both policies are the same. To the right of the solid curve, we have the parameter configurations for which an incumbent experiments in period 1.

Proposition 1. *A representative agent who serves two periods in the absence of elections adopts the risky decision in period 1 when* $pr > \frac{1+p}{2}$ *and so takes risks in the first period even when her expected one period pay off from doing so is less than that obtained when playing safe.*

Increasing term length increases an agent's willingness to take risks. She takes risks even when the expected first period payoff from doing so is *less* than obtained when playing safe. The intuition is immediate. An agent who chooses the risky policy in the first period and is successful will do so again in the second period. An agent who takes the risky policy and fails will play safe in the second period. In choosing the risky option the agent learn about her type and thereby improve her payoff. Thus it is optimal for her to engage in policy learning. The comparative statics for this example are straightforward: fixing the value of the risky option then an *ex-ante* more competent agent engages in more first period risk-taking. Fixing the competence of the agent, she engages in more risk-taking when the policy payoff conditional on a success is larger.

Figure 2 provides a graphical illustration. The 45° line defines $pr = 1$ where expected returns from safe and risky are equal and so above this line first period expected payoffs from choosing risky are lower. The solid curve separates the parameter space according to the incumbent's optimal actions: to its right she chooses risky, to the left she plays safe. When p is low then the politician may still takes risks in a 2 period model if r is large. Similarly when p is large then efficient risks may be taken for a larger set of values.

5. Policy Learning with Noncompetitive Election

We now introduce electoral competition: after the incumbent has chosen policy and the outcomes are revealed the electorate can choose to reelect her or elect a different politician to office. This is perhaps the most reduced form version of an election that we can consider: since we set $q_o = 0$ the electorate does not obtain information about the opponent. Then forward looking voters choose between alternatives based only upon the observed record of the incumbent.[7]

Proposition 2. *When the incumbent faces reelection then the subgame perfect equilibrium involves her adopting the risky policy in period 1 only when* $pr > \frac{3}{4}$ *and playing safe otherwise. When the incumbent plays risky she is reelected if successful and replaced otherwise. If she plays safe then is is reelected with probability* $1/2$*. Only a successful incumbent chooses the risky option in the second period, otherwise the elected politician plays safe in the second period.*

The main insight gleaned for a comparison of proposition 1 and 2 is that an election induces risk-taking behavior that is inefficient.

Corollary. *In a noncompetitive election, when* $p < 1/2$*, then selection of the incumbent induces first period investment in the risky option that is (weakly) lower than the optimal level, whereas when* $p > 1/2$ *selection induces the incumbent to (weakly) overinvest in the risky policy. In comparison to the appointed official, the introduction of a noncompetitive election leads to inefficiency.*

Although the voter is concerned with the efficient level of policy risk she is unable to induce such outcomes via selection of the incumbent in a noncompetitive environment. Indeed selection induces worse outcomes for the voter relative to the benchmark where the policymaker is appointed.

At low levels of competence the introduction of selection means that the politician under invests in the risky policy relative to the efficient first period benchmark. A politician fears that when taking the risky action the outcome will be a failure and thereby she will reveal herself to be incompetent. This "fear of failure" leads to an inefficiently low level of investment. The paradox is that when the expected competence of the incumbent is low, and so learning is most valuable to the voter, then a politician, whatever her type, will prefer to play safe.

At high levels of competence, by contrast, the politician over invests in the risky policy relative to the efficient benchmark. In these situations, although the average competence of the politician

[7]Selection is common to many decision-making environments. A distinction here is that the voter is concerned with efficiency not private interest.

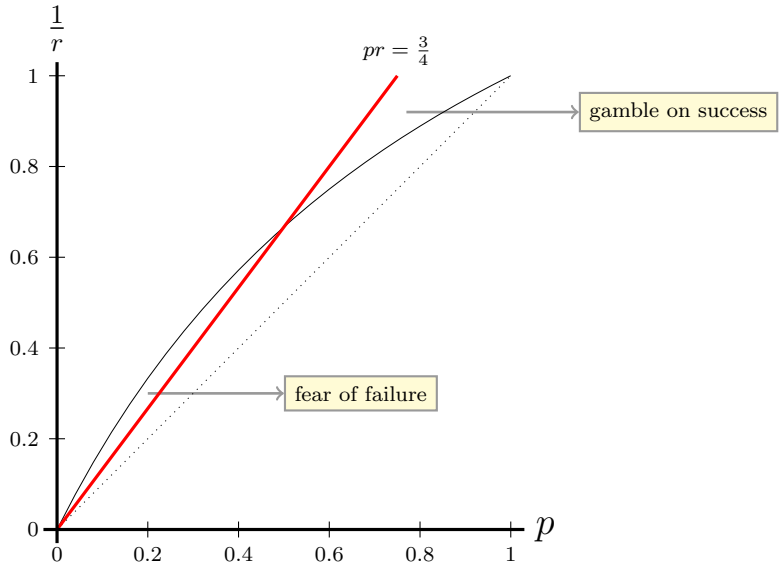

FIGURE 3. First Period Policy with Selection. The dotted line indicates $pr = 1$ where expected payoffs from both policies are the same. To the right of the thick solid line, we have the parameter configurations for which an incumbent experiments in period 1 (the differences between the solid curve and the thick solid line indicate that first period experimentation is weakly inefficient relative to the benchmark).

is high, the value of the risky option is relatively low. Then voters would rather the politician plays safe. The politician anticipates, however, that choosing the risky option will likely yield a success, that the electorate will observe this, and that she will be reelected. And so she "gambles on success".

Figure 3 provides an illustration of these inefficiencies induced by selection of the incumbent in a noncompetitive election. For $p < 1/2$ the politician's fear of failure leads to an inefficiently low investment in the risky policy. When $p > 1/2$ the politician gambles on success thereby adopting an inefficiently high level of investment in the risky policy. The conflict of interest between the politician and the voter arises due to the career concern of the former. The electoral environment that we consider here is the same as that considered in the literature on "pandering" (Canes-Wrone, Herron, and Shotts, 2001; Maskin and Tirole, 2004), though both the policy environment and informational structure are different. Indeed our central finding–that elections can produce inefficiencies due to office-seeking behavior by politicians–suggests the robustness of those earlier results to such differences even with symmetric information.

6. Policy Learning with Competitive Elections

We now model the actions of an incumbent who faces a competitive election. After choosing the first period policy, the public observes the policy outcome. In contrast to the previous case, however, in

a competitive election the opponent is now an active participant. She chooses her position before voters form their final opinion. We assume that the opponent can influence (though not control) the perception of the public. The uncertainty between the opponent's announcement and the voters evaluation is resolved in the interim before the election.

To keep the analysis simple we first assume a perfectly symmetric situation between the incumbent and the opposition by setting $q_j = 1$ for $j \in \{i, o\}$. This implies that during the course of the campaign the challenger can convince the electorate of her competence; indeed she can be as convincing *as if* she had successfully implemented the risky policy. We show that even in this extreme case–where the words (of the challenger) speak as loudly as the deeds (of the incumbent)– then under and over investment by the incumbent in the risky policy remains.

First we establish the subgame perfect equilibrium in a competitive environment.

Proposition 3. *Suppose that $q_j = 1$ for $j \in \{i, o\}$. In a perfectly competitive election the subgame perfect equilibrium involves the incumbent playing risky in period 1 when $pr > f(p)$ and safe otherwise. If the incumbent plays risky and fails, the opponent plays risky if $pr > \frac{1+p}{2}$ and safe otherwise. If the incumbent plays risky and succeeds, the opponent plays risky. If the incumbent plays safe then if $pr > 1/2$ the opponent plays risky whilst playing safe otherwise. An incumbent who plays risky is re-elected with probability $1 - \frac{p}{2}$ if successful; whilst an incumbent who plays safe is reelected with probability $1/2$ if her opponent does likewise and $1 - p$ otherwise.*

The proposition separates the parameter region according to the function $f(p)$ that is explicitly defined in the appendix and illustrated below. Before that illustration we first, and as before, use our equilibrium result to analyze policy learning in a competitive electoral environment relative to the efficient benchmark.

Corollary. *Suppose that $q_j = 1$ for $j \in \{i, o\}$. When p is low then an election with competition induces first period investment in the risky option by the incumbent that is (weakly) lower than the optimal level, whereas when p is high then an election with competition induces the politician to (weakly) overinvest in the risky policy.*

The result establishes an important finding: the introduction of an active opponent in a competitive election does not alter our key finding that holding the official to account by election induces inefficiency with respect to first period investment in the risky option by the incumbent.

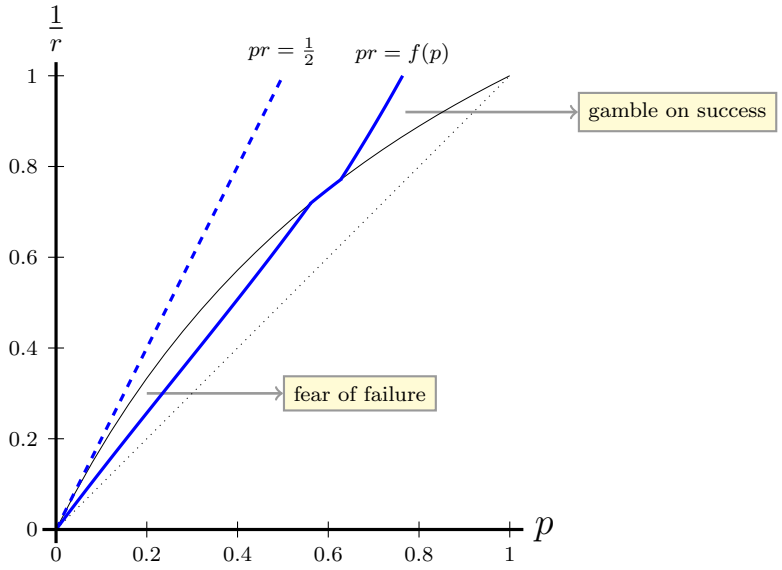

FIGURE 4. First Period Policy with Competition. The dotted line indicates $pr = 1$ where expected payoffs from both policies are the same. To the right of the thick solid line, we have the parameter configurations for which an incumbent experiments in period 1 (the differences between the solid curve and the thick solid line indicate that first period experimentation is weakly inefficient relative to the benchmark). To the right of the thick dashed line, we have the parameter configurations for which the opposition runs on a risky platform when the incumbent plays safe.

Figure 4 shows the incumbent's action relative to the efficient benchmark. At low levels of p first-period investment is lower; at higher levels of p it is higher; we also observe that for some intermediate range of competence values, investment by the politician who faces competitive elections is the same as that of an appointed policymaker. Our result and accompanying illustration simply reinforces our earlier observation: a fear of failure leads to first period investment in the risky policy that is lower than that taken by the appointed official at low levels of competence; whilst gambling on success arises at higher levels of competence. In fact we can go further in comparing first period risk taking across perfectly competitive and non competitive electoral environments

Proposition 4. *Risk-taking by incumbents in the first period is unambiguously lower under perfectly competitive elections than noncompetitive elections.*

The intuition for this finding is most immediate at high levels of competence where, as we have seen, the introduction of elections induces excessive risk-taking. Whilst a priori one would think that increasing the competitiveness of the election would exacerbate risk-taking behavior, in fact the opposite is true. Whereas previously the incumbent was always reelected following the successful implementation of the risky policy, now she may not reap the full reward from successful adoption of the risky platform. Instead, in a competitive election, her opponent can also reveal herself to

be competent. Then the incumbent and opponent are reelected with equal probability. This curbs excessive risk-taking by the former.

Proposition 4 also states that at low levels of competence risk-taking is lower under competitive elections. The intuition is somewhat more subtle. We have already observed that when competence is low on average then the incumbent exhibits fear of failure. Next note that when the incumbent plays safe her opponent can distinguish herself only when choosing the risky policy. Of course, given the symmetry in competence, for this range of the parameter space the opponent is more likely than not to be incompetent. Anticipating this reinforces the incentive of the incumbent to play safe. When her opponent chooses the risky policy she may reveal herself to be incompetent; but then in a perfectly competitive environment the incumbent is reelected with certainty. In this case, when playing safe, the incumbent gambles on her opponent's failure.

The underlying logic that links these scenarios is that competitive elections induce less risk-taking by incumbents and boldness by the opposition. Such behavior is sequentially rational and is observed in regular everyday politics: an opponent taunts the incumbent to take decisive action, and when none is forthcoming, (the incumbent sticks with tried and tested methods), accuses her of lacking leadership. A contemporary example involves the current Labour opposition in the United Kingdom (previously the government) who on several occasions have demanded bolder regulation of financial and media markets, whilst the current Conservative- Liberal Democrat government (previously the opposition) point to the failure of the now opposition to implement such measures when in office.

7. Welfare Analysis

It might appear from the preceding analysis that competitive elections reduces the level of learning via risk-taking by the incumbent and so enhance inefficiencies. We now show that this is not so.

Proposition 5. *When $pr < 1$ noncompetitive elections are (weakly) inefficient with respect to the benchmark in which the politician is appointed for two terms with no re-election. When $pr < 1$ competitive elections are (weakly) welfare improving with respect to noncompetitive elections and with respect to the benchmark only for low levels of competence. When $pr > 1$ both selection and competition enhance efficiency.*

Comparing across cases we find that appointment of the politicians is (for $pr < 1$) at least as good and sometimes preferable to noncompetitive elections. However, for a large range of the

parameter space competitive elections are preferable with respect to both noncompetitive elections and appointment.

The efficiency gains under noncompetitive elections are due to the best response profile of the incumbent and her opponent. As we have also observed, the introduction of competitive elections with an active opponent induces the incumbent to take fewer risks than she would do otherwise. Moreover, and as we have also seen by Proposition 3, when the incumbent chooses the safe option her opponent may choose the risky one. This is the source the efficiency gain.

To see this note that the electorate evaluates the expected benefit from implementing the opposition's risky policy against the incumbent's safe option. But before casting her vote the voter learns about the risky option. More precisely, she learns whether the imposition of the risky policy by the opponent will yield a successful outcome. Indeed in the case of perfect competition the opponents type is revealed when she campaigns on the risky platform. Critically, the voter learns these facts without incurring the costs of policy policy implementation (ie. the forgone payoff from the safe option when the risky policy yields a failure). The cost, if there is one, is born by the opponent. If she fails when choosing the risky policy then she loses the chance of holding office. But this does not have a negative effect on voter welfare.

As we will illustrate below, across the range of the parameter space where the prior competence of the incumbent is low–and so policy failures are on average more likely should the incumbent choose the risky option–the benefit to the voter of having a strategically rational opponent is high.

Figure 5 illustrates the key regions for comparison: the left hand figure shows outcomes under noncompetitive elections; the right hand figure shows outcomes under competitive ones with an active challenger where we set $q_o = q_i = 1$. The shading in these figures depicts the comparison with the benchmark case whereby the incumbent is appointed: dark (red) shading depicts inefficiency relative to the benchmark; light (green) shading indicates superior voter welfare relative to the benchmark appointment case; whereas no shading indicates equivalence in welfare terms. The figure clearly shows that noncompetitive elections weakly reduce efficiency with respect to the benchmark whenever $pr < 1$; by contrast, and for a large range of the parameter space, competitive elections are welfare enhancing.

When $pr > 1$, both competitive and non-competitive elections are strictly welfare improving with respect to the benchmark case. This is so because an appointed policymaker cannot be replaced by voters even when she is not competent. Such a politician will always be replaced under both types

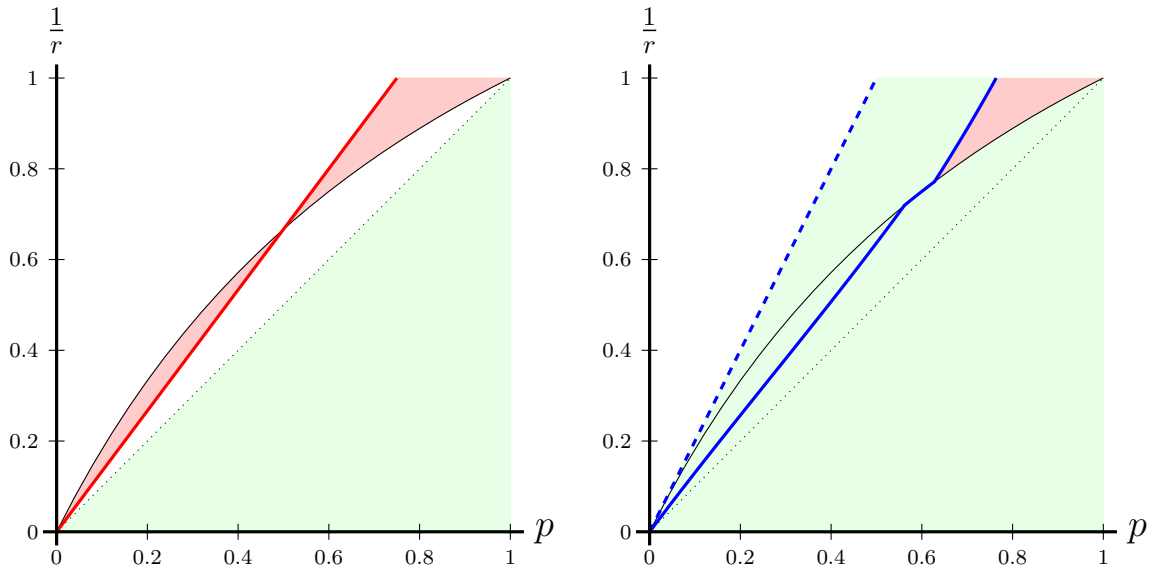

FIGURE 5. Competitive-vs-Noncompetitive election. The left hand figure shows outcomes where $q_o = 0$ and $q_i = 1$ and so the incumbent does not face an active opponent. The right hand figure shows outcomes with perfect competition $q_o = q_i = 1$. Dark (red) shading indicates inefficiency relative to appointment of the politician. Light (green) shading indicates strictly superior outcomes than appointment. No shading indicates equivalent outcomes as under appointment.

of elections. In the case of noncompetitive elections the efficiency gains with respect to appointment are isolated for $pr > 1$.

Note that the light (green) shaded area in the parameter range $pr < 1$ indicates precisely the parameter range for which competitive elections outperform noncompetitive ones. To the right of the dashed thick (blue) line lies the area where the opponent chooses the risky policy when the incumbent plays safe. To the right of the solid blue line, the area where the incumbent chooses the risky policy. The shaded area between these lines depicts the area where the incumbent plays safe whilst her opponent takes risks. This is the parameter region of costless learning where some of the welfare gains are located. There are also welfare gains to the right of the solid blue line where *both* incumbent and opponent choose the risky policy in a competitive elections and so the voter always learns more relative to the other cases of comparison. Of course, in Figure 5 we only compare perfectly competitive ($q_o = 1$) elections with non competitive elections ones ($q_o = 0$). Do our insights about the benefits of competitive elections go further? We show the robustness of our insights for cases where $q_o \in (0, 1)$.

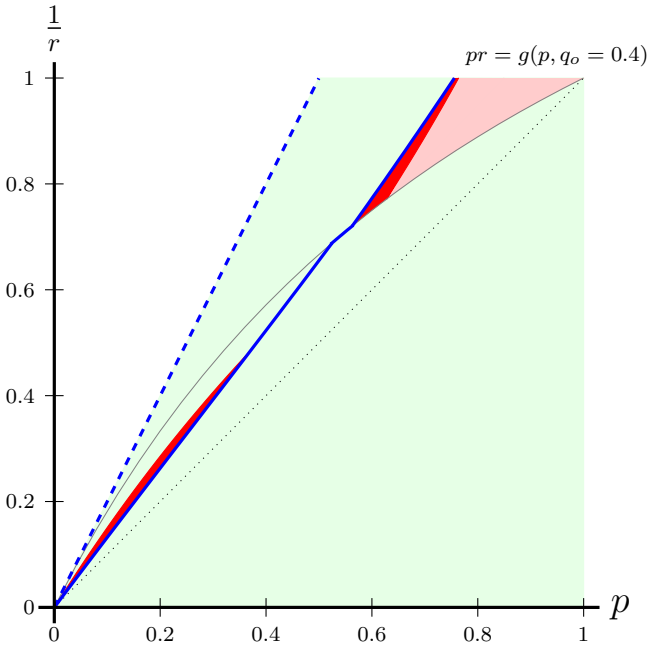

FIGURE 6. Competitive Election with Imperfect Competition ($q_o = 0.4$). The dotted line indicates $pr = 1$ where expected payoffs from both policies are the same. To the right of the thick solid line, we have the parameter configurations for which an incumbent experiments in period 1. To the right of the thick dashed line, we have the parameter configurations for which the opposition runs on a risky platform when the incumbent plays safe. Dark (red) shading indicates inefficiency relative to appointment of the politician; the very dark (red) shading indicates the increased inefficiency with respect to the competitive case when $q_o = 1$. Light (green) shading indicates strictly superior outcomes than appointment. No shading indicates equivalent outcomes as under appointment.

Proposition 6. *An increase in the competitiveness of the election leads to a decrease in first period risk-taking by the incumbent: the sub-game perfect equilibrium involves the incumbent playing risky in period 1 when* $pr > g(p, q_o)$*, where* $\frac{\partial g(p,q_o)}{\partial q_o} \leq 0$*. This implies that the inefficiency owing to the incumbent 'gambling on success' (high values of p) or that are due to her 'fear of failure' (low values of p) strictly decrease with* q_0*.*

Earlier we contrasted the case where $q_o = 0$ with that of $q_o = 1$ and showed that the level of first period risk-taking was lower under the latter fully competitive elections. Here we show that this effect is continuous and monotonic in q_o. As the competitiveness of the election increases the level of risk taking by the incumbent decreases.

We explore the welfare effect. The second part of proposition 6 reveals that welfare always increases in the competitiveness of the election. This is illustrated in Figure 6 that contrasts a semi-competitive election ($q_o = 0.4$) with a fully-competitive one ($q_o = 1$). The area depicting

inefficiencies that are due to an incumbent's gambling on success in a fully competitive election are now larger with restricted competition. Moreover, whereas inefficiencies that occur when $p < 1/2$ are eradicated in a fully competitive election, they reappear in a partially competitive one. The reason for this is that although, and as argued above, voters learn from advocacy of the risky policy by the opponent, they do not learn as much as if the incumbent actually implemented the risky policy. The inefficiencies that arise when moving from a fully competitive to a partially competitive one are illustrated in Figure 6 in the darker (red) shaded areas of the relevant (inefficiency) region.

A clear intuition emerges. Sufficiently competitive elections allow voters to learn policy relevant information without incurring costs of experimentation. This speaks to the desirability of elections as control mechanisms. However, and although Proposition 5 shows that elections can lead to efficiency gains over and above those that could be attained by an appointed official, some inefficiency remains. In particular, as we have shown, elections can lead to overzealous adoption of the risky project when doing so is not in the public interest. These effects are particularly strong under noncompetitive elections. Then the incumbent will always be re-elected if she successfully implements the risky option and so goes for glory despite the consequences. They are still present, however, in competitive elections that, whilst ameliorating the effect, do not eradicate it entirely.

The reader might suspect however that under an optimally designed scheme the voter could do better. For the voter might punish an incumbent who boldly experiments when it is not in the public interest that she do so, or punish an incumbent that shies away from risky policies despite public demand for such measures. The question remains whether she can credibly do so. We turn to an analysis of this question.

8. Selection and Incentives

The question we ask is whether voters can simultaneously use their votes to select high performers and provide incentives for better performance such that any remaining inefficiencies illustrated earlier can be removed. Logic suggests otherwise. Once the incumbent has gambled on success and proven herself to be competent, it is rational for the voter to retain him in office. Equally, when a politician fails then the voter may be unable to pre-commit to retaining him in office. This logic follows closely the argument made by Fearon in his well known critique of elections as incentive mechanisms (see also Besley (2006)). There he argued that voters *could not* simultaneously use their votes to select high performers and provide incentives for better performance. This logic, whilst compelling, is however flawed with respect to our policy environment.

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 and competitive elections.*

An immediate implication follows, as illustrated below in Figure 7:

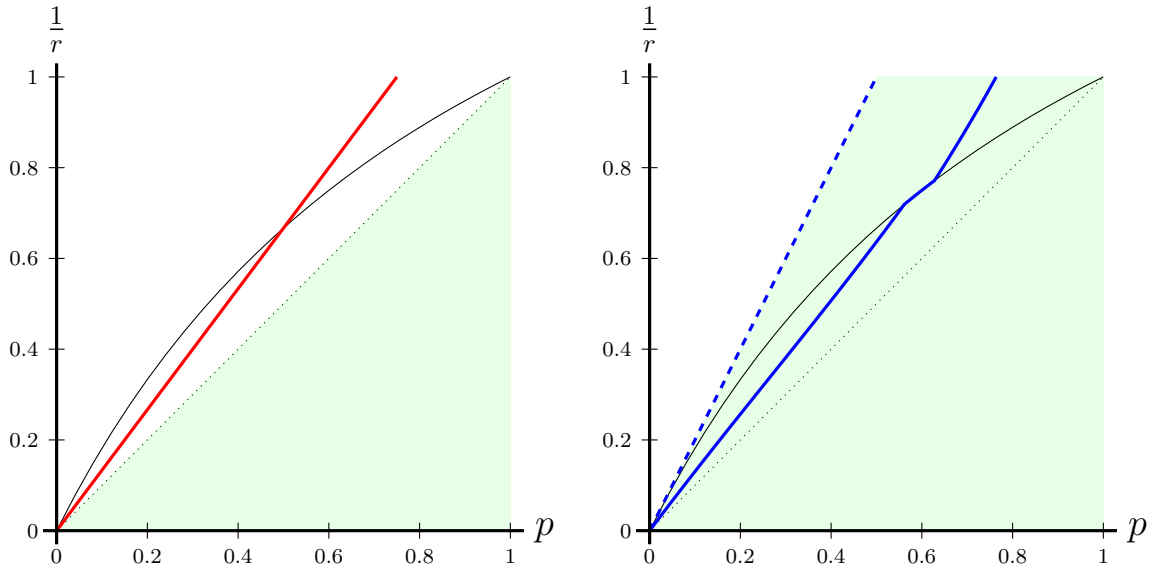

FIGURE 7. Incentives Recover Efficiency. The left hand figure shows outcomes with non-competitive elections. The right hand figure shows outcomes with competition. Light (green) shading indicates strictly superior outcomes than appointment. No shading indicates equivalent outcomes as under appointment.

Proposition 8. *When $pr > 1$ both competitive and 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; instead, competitive elections are strictly welfare improving with respect to the benchmark whenever $pr > 1/2$ and equivalent when $pr < 1/2$.*

We relegate the details of our argument to the appendix. Here we provide some intuition. Under noncompetitive elections, voters can recover the benchmark level of welfare by rewarding the incumbent with reelection when she does not "gamble on success". They can reelect him with probability one when she implements the safe action. Moreover they can punish the incumbent by not selecting him when she under invests in the in the risky policy. Thus when voters use an optimal incentive scheme in a noncompetitive election then inefficiency is alleviated as illustrated in the left-hand side panel of Figure 7.

Under competitive elections, voters can no longer use the same incentive scheme. When $pr > 1/2$ and the incumbent plays safe, then the opponent will choose the risky platform. (This is the area bounded by the dashed blue line on one side and the solid blue line on the other on the right hand side panel of Figure 7). Then the voter cannot credibly reward the incumbent for her restraint as her evaluation of the opposition platform determines the outcome of the election. In particular, if the opponent is successful then she always will be elected.

This does not mean, however, that the voter can not provide incentives. Note that when both incumbent and opponent choose the risky policy and are successful, a sequentially rational voter is strictly indifferent between retention or replacement of the incumbent. The voter can then credibly commit to breaking the tie *in the opponent's* favor. That is, she selects the opponent with probability one should both the incumbent and the opposition adopt the risky policy and be successful. This simple tie-breaking rule is enough to provide the correct incentives to the incumbent who now pursues the efficient course of action.

Before completing this section we provide a full analysis of the equilibrium in competitive elections.

Proposition 9. *When voters use their votes to both select incumbents and to sanction their performance then in a competitive election the subgame perfect equilibrium involves: the incumbent and the opponent playing safe if* $pr < 1/2$*; if* $1/2 < pr < h(p)$ *then the incumbent plays safe whilst the opponent plays risky; finally if* $h(p) < pr$ *then both the incumbent and her opponent play risky.*[8] *When only one player (either incumbent or opponent) plays risky then she is (re)-elected only if successful and if unsuccessful is elected with probability zero. If both players choose the risky option and are successful then if* $f(p) < pr < h(p)$ *the opponent is elected whereas if* $h(p) < pr < f(p)$ *the incumbent is elected; for all other values of pr they both have an equal chance of being elected. Finally, when both play risky and only one is successful this one is elected and when they are both unsuccessful, they both have an equal chance of being elected.*

The equilibrium describes fully the optimal incentive scheme for the voter. It involves comparison of two possible scenarios: in the first, the voter does not incur the costs of policy experimentation but benefits when the opponent takes risks; in the second, the voter wants to induce as much policy learning as possible and can achieve this by ensuring that the opposition chooses the risky policy even when the incumbent has done so and failed. Proposition 8 establishes the parameter values for

[8] h(p) is explicitly defined in the appendix and illustrated in figure 8 below.

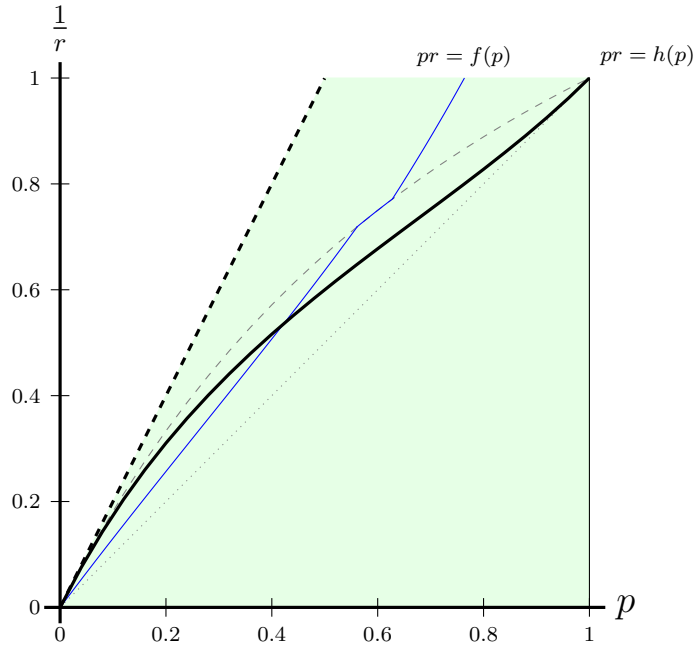

FIGURE 8. Equilibrium with Selection and Incentives. To the right of the thick solid line, we have the parameter configurations for which both the incumbent and the opposition experiment in period 1. To the right of the thick dashed line, we have the parameter configurations for which the opposition runs on a risky platform when the incumbent plays safe. Light (green) shading indicates strictly superior outcomes than appointment. No shading indicates equivalent outcomes as under appointment.

which each of the two situations is superior. The accompanying depiction (Figure 8) illustrates that over some range of the parameter space the voters induce risk-taking only by the opponent thus curtailing the incumbent's desire to gamble on success. Over a different range the optimal incentive scheme induces both opponent and incumbent to take risks, thus curtailing their respective fear of failure. The thick solid (red) curve in Figure 8 defines $h(p)$ that separates these two regions: to its left the opponent only takes risk; to its right both the incumbent and her opponent take risks.

We conclude this section by highlighting the implications of our analysis for the study of incumbency advantage. Existing studies leave open whether this advantage is systemic or whether it arises due to the actions taken by an incumbent that make her preferable to voters. For example, the partisan incumbency advantage identified in several studies may be due to systemic bias, voters' preference for electing established candidates, or due to the policy positions adopted by incumbents. Our theoretical framework allows us to analyze optimal voting when both candidates are identical ex-post and so voters can not have preferences for one over the other.

As we have seen when the incumbent is perceived to be of high quality and gambles on success, the voter can credibly commit to punishing her when (ex-post) indifferent between the candidates. In

this case the incumbent who stands on a record of successful implementation of the risky policy will not be rewarded if her opponent, who likewise stands on the risky policy is able to reveal herself to be competent. This is a source of incumbency disadvantage.[9] In this case the voter systematically biases against the incumbent though this bias is not exercised on the equilibrium path.

However, as we have also seen, the voter can credibly commit to rewarding an incumbent of perceived low quality who takes risks. Then if both incumbent and opponent adopt the risky policy and are successful, the voter commits to reelecting the former. This is a source of incumbency advantage. Voters reward incumbents for establishing a policy record and this systematic bias is exercised on the equilibrium path.

We note two interesting features with respect to competitive elections. The first is that under the voters optimal incentive scheme the platform of incumbent and challenger diverge. In particular, to the left of the thick solid curve for $h(p)$ indicated in Figure 8 the incumbent plays safe whilst her opponent plays risky. The second insight is that a commonly perceived source of incumbency advantage, namely inherent competence, can work against the incumbent. Here incumbents who are on average of high competence and may reveal this by taking inefficient risk are punished (of the equilibrium path) in the optimal scheme if the opponent can reveal herself to be equally competent. This insight complements that of Gordon, Huber, and Landa (2007) and Gordon and Landa (2009) who show that under symmetric information costly entry by a challenger in a race where the incumbent is perceived to be of high competence can counteract this source of incumbency advantage.

To complete our analysis we look at incumbency advantage/disadvantage in the absence of an active challenger. Then straightforwardly, as we have seen, incentives for better incumbent performance require that she be rewarded when refraining from unwarranted risks and punished when shying away from such risk. When the prior for competence is high, incumbents are rewarded for playing safe when it is efficient that they do so: this is an incumbency advantage. But they are punished when the prior is low and they shy away from implementing the risky policy when it is inefficient they do so: this is an incumbency disadvantage.

We record the insights in this section as a corollary to our earlier result.

Corollary. *In a competitive election the incumbent is disadvantaged when a priori she is of high competence and advantaged otherwise. In a noncompetitive election the incumbent is advantaged when a priori she is of high competence and disadvantaged when she is of low competence.*

[9]The voter cannot credibly commit to reward an incumbent that plays safe in a competitive election.

9. An Experiment on the Optimal Level of First Period Experimentation

Our model is of risk and learning by politicians in competitive and noncompetitive environments and its analysis leads to testable hypothesis with regard to the first period actions taken by the incumbent. The first follows immediately from Proposition 1: in the absence of elections, an agent will adopt the risky policy in the first period even when the expected one period payoff is less than that obtained when playing safe. This reveals that an agent is willing to take risks in order to learn. The second hypothesis follows from the key (corollaries to) Propositions 2 and 3 which state that, relative to the efficient levels of risk taken by an agent appointed to two consecutive terms, a politician who faces an election after the first period will take fewer risks when her expected competence is low, she "fears failure". The third hypothesis follows from the same corollaries which also state that, relative to the efficient levels of risk taken by an agent appointed to two consecutive terms, a politician who faces an election after the first period will take higher risks when her expected competence is high, she "gambles on success".

Our key results are related to the discrete "up" or "down" nature of the winner-take-all contest: once a politician has revealed herself to be competent she will be rewarded by the electorate; likewise she will be punished when she proves herself to be incompetent. Although, as we show, the voter can use sanctioning in combination with selection to achieve efficient outcomes, it is interesting to explore whether (in the absence of such sophisticated voter strategies) the risk profile described by our model will arise. As noted already in our introductory remarks, the fact that our twin hypothesis of "fear of failure" and "gambling on success" that relate to the role of elections run counter to predicted behavior under prospect theory (Kahneman and Tversky, 1979). This suggests that these effects can be identified by empirical study in a suitably controlled environment.

To explore this we design an experiment that replicates the conditions under which incumbents make their first period choices under a winner-take-all election. 123 subjects took part in our experiment that was implemented at the Center for Experimental Social Sciences at Nuffield College, Oxford. Subjects participated in a sequence of lotteries in which they could either obtain a sure payoff (the safe option) or a risky payoff the values of which depended on two parameters: p the probability that they secured the prize and r the expected value of the prize. Subjects faced 8 combinations of parameters (depicted in figure 9) and for each parameter configuration they had to decide between a safe or a risky option in four different treatments representing a one period model (1P), a two period model with appointment (2P), a two period model with non-competitive election after period 1 (NCE), and a two period model with competitive election after period 1 (CE). The treatments

reflect expected payoffs given the anticipation of sequentially rational behavior by voters (under NCE) and voters and an opponent (under CE). Full details of the experimental design are provided in Appendix B.

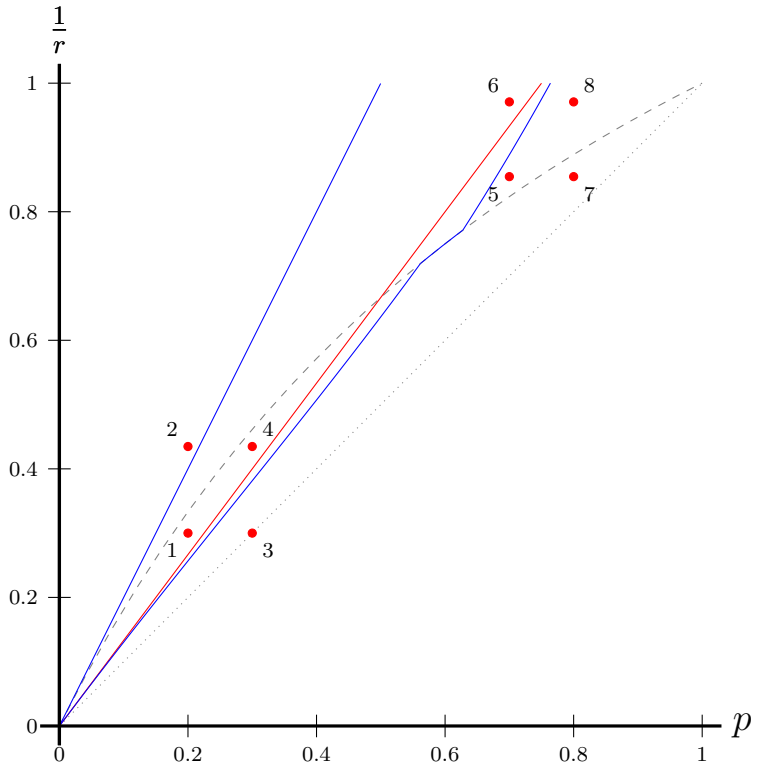

FIGURE 9. Parameter Combinations Used in Laboratory Experiment. Each point in the graph corresponds to a combination of parameters: competence p and return of risky option r.

Given our parameter combinations illustrated in Figure 9 we can unequivocally state our hypothesis: for combinations 1, 3, 4 and 7, we expect to observe subjects adopt the risky choice more frequently when confronted with a situation representing a two period model than with a situation representing a one period model (hypothesis H1); for combinations 1 and 4, we expect to observe fear of failure by which our subjects should be choosing the risky option less often under any sort of election than under appointment (hypothesis H2); instead, for combinations 5 and 8, we expect to observe gambling on success with subjects under the competitive and non-competitive elections choosing the risky option more frequently than under appointment (hypothesis H3). Table 1 summarizes our hypothesis with respect to the subjects' anticipated behavior under each of the treatments.

TABLE 1. First Period Investment in Risky Policy under Different Treatments and combinations of parameters

Hypothesis	Treatments	Parameter Combination	Directional Prediction
H-1	2P and 1P	1, 3, 4 and 7	$2P > 1P$
H-2	NCE and 2P	1 and 4	$NCE, CE < 2P$
H-3	CE and 2P	5 and 8	$NCE, CE > 2P$

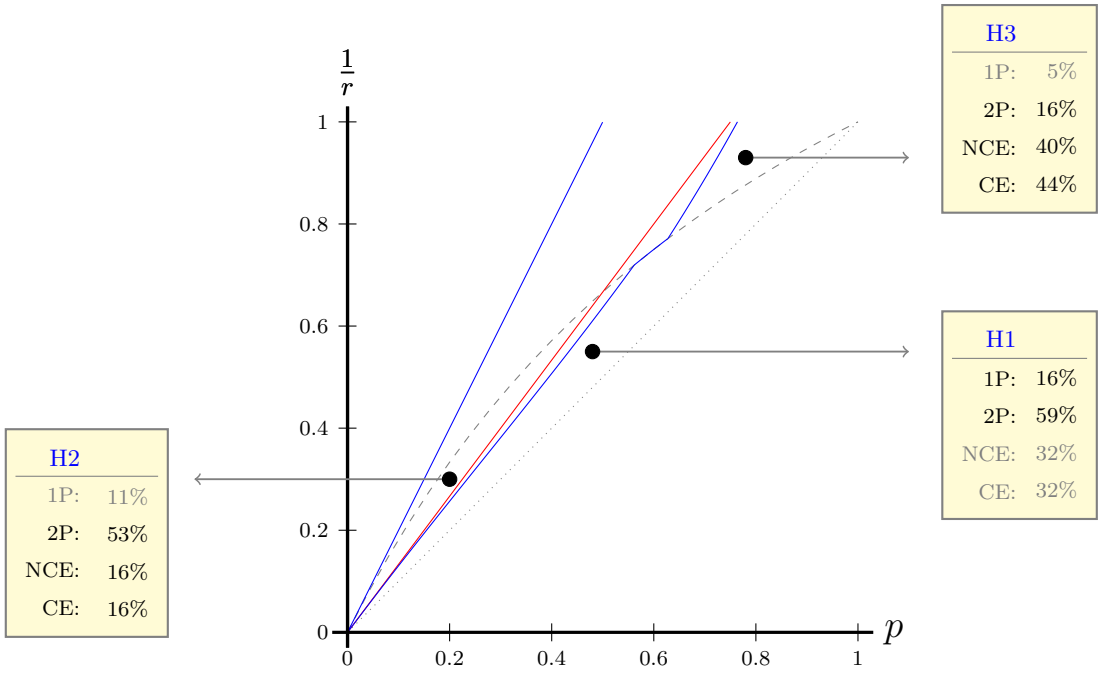

FIGURE 10. The percentage of subjects choosing the risky option in each of our three predicted hypothesis for our 4 treatments representing a one period model (1P), a two periods model (2P), a noncompetitive election (NCE), and a competitive one (CE).

Figure 10 provides a summary of our key findings: under (1P) we show the percentage of subjects who chose the risky option when faced with a choice representing a one period only; under (2P) the percentage of subjects who chose the risky option when faced with a choice representing two periods; under NCE we record the choice of the risky alternative when subjects faced the same two period choice with the payoffs as given when rational voters select on first period performance; finally (in CE) we record the outcomes with expected payoffs as in NCE but adjusted to incorporate the rational actions taken by voters in a competitive election involving a (rational) opponent. For each hypothesis we indicate in bold the relevant comparisons. (In gray we denote the cases that are not relevant for each of our three hypothesis).

In directly testing our hypothesis we first compare outcomes under 1P and 2P for each of these parameter configurations.[10] Doing so we observe that our first hypothesis, that agents will engage in learning, is supported by the data. In (H1) 59% were willing to choose the risky option under treatment 1P whereas under 2P 16% were willing to do so.[11]

Our second hypothesis (H2), that elections induce fear of failure at lower levels of competence is observed by comparing outcomes under treatment 2P and treatments NCE and CE. If our hypothesis is correct then relative to the outcomes under 2P we should observe a reduction in first period risk taking under NCE and CE. Indeed we observe that almost all of the increase in risk taking that occurs when subjects faced 2P rather than 1P is eradicated when payoffs reflected retrospective evaluation in treatment NCE and CE.

Our third hypothesis (H3), that that elections induce gambling on success at higher levels of competence is also observed by comparing outcomes under treatment 2P and treatments NCE and CE and is also supported by the data. We observe that risk taking increased dramatically under the same incentives when levels of competence was high.

In sum, Figure 10 presents solid evidence in favor of each of our three hypothesis.

Moving beyond this pictorial summary, in Table 2 we provide average levels of investment in the risky option for each of our treatment conditions and parameter configurations illustrated in Figure 9. It is straightforward to confirm that each of our directional hypothesis listed in Table 1 hold in the data. In particular investment in the risky option is higher under treatment 2P than treatment 1P under all combinations (learning through experimentation); investment is lower under treatment NCE and CE than under 2P for combinations 1-4 (fear of failure) and higher for combinations 5-8 (gambling on success). Note however that our hypothesis hold more strongly for the combinations of parameters where our theory predicts differences. For example the differences between 2P and 1P are largest in combinations 1, 3, 4, and 7.

Each subject faced all four treatments and all parameter combinations. Moreover, our hypothesis is directional so a simple one-sided paired difference in means test suffices to establish the statistical significance, if any, of our findings. Given the magnitude of the difference in the estimates it is no surprise that we can reject the null hypothesis that the difference in means for H1, H2 and H3 are zero at the 99% confidence level. Table 3 provides the difference in means for each of the relevant

[10] In evaluating our results we first note that, in all cases, some subjects were willing to take inefficient risks when faced with a single period choice when it was better they chose the safe option.

[11] Note that similar comparison holds in (H2) because the set of parameter combinations of (H2) is a subset of set in (H1).

TABLE 2. Average First Period Investment in Risky Policy under Different Treatments

		1P	2P	NCE	CE
Combination of Parameters	**1**	0.172	0.570	0.205	0.184
	2	0.045	0.381	0.090	0.094
	3	0.246	0.779	0.373	0.369
	4	0.053	0.500	0.123	0.135
	5	0.086	0.246	0.410	0.512
	6	0.008	0.070	0.164	0.193
	7	0.160	0.500	0.574	0.611
	8	0.008	0.082	0.389	0.361

comparisons. Since our predicted behavior is contrary to that predicted in prospect theory we can be sure that the effects are not confounded by the behavioral aspects highlighted by Kahneman and Tversky (1979). Moreover since such behavioral aspects would affect each of our treatments they would likely bias our results downwards.

In sum, the data from our experiment supports the notion that an incumbent who is held to account for her first period actions will react differently to one who is appointed for two terms. Although our model anticipates more sophisticated behavior than we have allowed for here, we nevertheless find it encouraging that some of our main hypothesis are supported out by the data. Here we have analyzed first period investment in risky policies by subjects (incumbents) given a set of optimal best responses from opposition politicians and voters. Extensions of the experimental design adopted here should assess whether in fact subjects do play the anticipated best response profiles. Nevertheless our model has passed its first test. The data shows evidence of the twin effects–fear of failure at lower levels of competence and gambling on success at higher levels of competence–that drive our results and arise naturally due to the distinct winner-take-all aspect of plurality rule elections.

10. CONCLUSIONS

The classic model of electoral accountability has that a politician will anticipate the reaction of voters to her performance at the time of election and that this induces better performance. Recent work, however, challenges this view. A growing literature suggests that politicians who are better informed as to the best course of action will instead pander to public prejudice thus leading to inefficient policies relative to those that would be implemented if the policymaker were appointed. We model an environment in which voters retrospectively evaluate politicians according to their performance but politicians are unsure as to the best course of action to take. Policy outcomes

TABLE 3. First Period Investment in Risky Policy Hypothesis Tests

Difference in Means		**1P and 2P**	**2P and NCE**	**2 and CE**
Combination	**1**	0.398	-0.365	-0.385
of Parameters		$(0.040)^{H1}$	$(0.041)^{H2}$	$(0.040)^{H2}$
	2	0.336	-0.291	-0.287
		(0.034)	(0.036)	(0.036)
	3	0.533	-0.406	-0.410
		$(0.038)^{H1}$	(0.041)	(0.041)
	4	0.447	-0.377	-0.365
		$(0.035)^{H1}$	$(0.038)^{H2}$	$(0.039)^{H2}$
	5	0.160	0.164	0.266
		(0.033)	$(0.042)^{H3}$	$(0.042)^{H3}$
	6	0.062	0.094	0.123
		(0.017)	(0.029)	(0.030)
	7	0.340	0.074	0.111
		$(0.040)^{H1}$	(0.045)	(0.045)
	8	0.074	0.307	0.279
		(0.019)	$(0.036)^{H3}$	$(0.035)^{H3}$

[a] We report difference in means for comparison treatments with standard deviation in brackets. All one-sided paired difference in means test statistics are significant at p<0.01 except that for H2 combination 7 for which $p(T < t) = 0.051$. With H1, H2 and H3 we denote the relevant tests for each of our three hypothesis.

reflect both the choice of policy and a politicians' ability; the politician is unsure of the latter. By implementing risky policies rather than safe ones a politician (and voters) can learn her type. We have shown that the standard model of electoral competition yields inefficient outcomes under these circumstances: politicians under invest in experimentation when they are a priori of low competence, whilst over-investing when of high competence. These outcomes can be explained as

owing to the particularly stark nature of the electoral reward. When voters select on performance then they reward success even when risk-taking was unwarranted, whilst punishing failure even when risks were justified. This leads to an excessive "fear of failure" at lower levels of competence, and encourages politicians to "gamble on success" when they are competent.

This inefficiency result is not robust, however, to a more realistic modeling framework whereby the incumbent faces an opponent in a competitive election. The distinguishing feature of an election as an accountability mechanism is precisely this competitive aspect. In no other profession are incumbents' actions scrutinized and evaluated relative to the pronouncements of an opponent in a winner-take-all contest. Indeed we find that the introduction of competitive elections ameliorate inefficiencies to a large extent. Moreover when voters use their votes to both select and sanction politicians then all efficiencies are eradicated and the performance of politicians elected under competitive elections is (weakly) superior to that of appointed officials. Our result that sufficiently competitive elections allow voters to learn policy relevant information without incurring costs provides a strategic rational and normative defence for constructive opposition. Our results chime with the views of Eldersveld who, in advocating two-party competition, stated that "the fundamental requirement of electoral accountability is a two-party system in which the opposition party acts as the critic of the party in power, developing, defining and presenting the policy alternatives that are necessary for a true choice in reaching public decision." Here we focus on the interaction between an incumbent and an opponent. Next we plan to explore the team elements of this problem.

Our model produces clear hypothesis concerning the expected level of investment in risky policies by an incumbent and we have subjected to empirical investigation. A laboratory experiment designed to capture incumbents incentives under different scenarios produces robust support for some of our main hypotheses. In particular we replicated the conditions under which incumbents made their policy choices under different scenarios in a winner-take-all contest. Payoffs reflected the optimal strategic behavior of opponents and voters. The data showed clear evidence of the twin effects–fear of failure at lower levels of competence, and gambling on success at higher levels of competence that arise naturally under a winner-take-all election.

11. Appendix A: Formal Proofs

Proof of Proposition 1. In the last period, an incumbent plays risky only when she knows that she is competent. This implies that the expected utility of period 1's incumbent from playing safe in period 1 is 2. If she plays risky then with probability p she obtains r in both periods and with

probability $1-p$ she obtains zero in the first period before and the sure payoff 1 in the second. It follows that in period 1 an incumbent implements a risky choice only when $pr > \frac{1+p}{2}$. □

Proof of Proposition 2. We find the subgame perfect Nash equilibrium by backwards induction. The voter always retains an incumbent who reveals herself as competent whilst replacing one who reveals herself as incompetent; otherwise (if the incumbent plays safe) she retains or replaces the incumbent with equal probability. In the second period, an incumbent plays risky if and only only she knows that she is competent. If she plays risky in the first period then with probability p she obtains r in both periods and with probability $1-p$ she obtains zero in both periods. Whilst if she plays safe then her expected payoff is $\left(1+\frac{1}{2}\right)$. It follows that in period 1 an incumbent implements a risky choice only when $pr > \frac{3}{4}$. □

Proof of Proposition 3. We find the subgame perfect Nash equilibrium by backwards induction. The voter re-elects a competent incumbent when her opponent reveals herself to be incompetent or plays safe; if, however, the opponent also reveals herself as competent then the incumbent is elected with probability 1/2. If the incumbent chooses the safe policy then she is elected with probability 1 if her opponent reveals herself to be incompetent; with probability 1/2 if the opponent plays safe; and zero probability if the opponent reveals herself as competent. Anticipating this voter behavior then the best response of the opposition to the first period choice of the incumbent is as follows: (1) when the incumbent chooses risky and is successful, the opponent has a positive probability of being elected and so receives a positive payoff if she also chooses risky and thereby is successful; (2) when the incumbent plays risky and is unsuccessful the opponent's expected payoff from playing risky is $pr + \frac{1-p}{2}$ whilst playing safe yields 1 and so she chooses risky if and only if $pr > \frac{1+p}{2}$; finally (3) when the incumbent plays safe, the opponent plays risky if and only if $(pr+(1-p)0) > \frac{1}{2} \Leftrightarrow pr > \frac{1}{2}$. Finally we solve for the incumbent's optimal first period choice when anticipating the opposition and voter's best response.

There are two cases to solve for. When $pr > \frac{1+p}{2}$ then the opposition adopts the risky policy even when the incumbent is unsuccessful. Anticipating this the incumbent adopts a risky policy when:

$$p\left[r+(1-p)r+\frac{pr}{2}\right]+(1-p)\left[(0+(1-p)\frac{1}{2}\right] > 1+(1-p) \tag{1}$$

The LHS of the inequality is the expected payoff from adopting the risky choice: the first term is the probability that incumbent is competent times her payoff in period 1 (r) plus her expected payoff in period 2 taking into account the possibility that the opponent may be successful and

then win the election with probability $\frac{1}{2}$; the second term is the probability she is not competent times her payoff in period 1 (0) plus her payoff in period 2 (she is reelected with probability $\frac{1}{2}$ only when the opponent is unsuccessful). Finally the RHS of the inequality is the expected payoff from playing safe in period 1. Solving inequality (1) we see that when $pr > \frac{1+p}{2}$, an incumbent adopts the risky policy only when $pr > \frac{3-p^2}{4-p}$.

When $pr < \frac{1+p}{2}$ the opposition adopts the risky policy only when the incumbent is successful. Then we obtain the same payoff from adopting the risky choice as in (1) with one difference, namely that the second term in the LHS of the inequality is now 0. The payoff from playing safe is $(1+(1-p))$ when the opponent's best response is to choose risky and $(1+\frac{1}{2})$ otherwise. The latter case applies when $pr > \frac{1}{2}$ and for such parameter values the incumbent always prefers to play safe. Instead, the former case applies when $pr < \frac{1}{2}$ and in these circumstances the incumbent adopts a risky policy only when $pr > \frac{4-2p}{4-p}$.

We thus conclude that the incumbent adopts a risky policy when $pr > f(p)$ where

$$f(p) = \begin{cases} pr > \frac{3-p^2}{4-p} & \text{if } pr > \frac{1+p}{2} \\ pr > \frac{4-2p}{4-p} & \text{if } pr < \frac{1+p}{2} \end{cases}$$

where the function $f(p)$ is illustrated in the main text. □

Proof of Proposition 4. To prove our result we need to show that the parameter values for which there is experimentation in perfectly competitive elections is strictly included in the set of parameters for which there is experimentation in non competitive ones. Recall from Proposition 2 (and Figure 5) that experimentation occurs under non competitive elections when $pr > 3/4$. Similarly, from Proposition 3 (and Figure 6) we know that experimentation occurs under perfectly competitive elections when $pr > f(p)$. There is strictly less risk-taking under competitive elections than under non-competitive elections if $f(p)$ is strictly larger than 3/4. Note that when $pr > \frac{1+p}{2}$ experimentation always occurs under non-competitive elections so the previous condition is satisfied. Instead, when $pr < \frac{1+p}{2}$ it can easily be shown that $frac4-2p4-p > \frac{3}{4}$ whenever $r \geq 1$. □

Proof of Proposition 5. The first claim, that when $pr < 1$ non competitive elections are inefficient with respect to the benchmark, is already established by the proof of proposition 3. Turning to the second claim, namely that when $pr < 1$ competitive elections are weakly welfare improving it suffices to show that the introduction of an active opponent does not induce any inefficiencies that are not present under noncompetitive elections and eradicates others that are. When $pr < 1$, as

shown in proposition 2, a non-competitive election leads to inefficiencies relative to the benchmark case of appointment that are due to over investment in the risky policy (for high p) and under investment (for low p). As shown in the proof of proposition 4 the first source of inefficiency is reduced in a competitive election. In a competitive election when $1/2 < pr < f(p)$, the incumbent plays safe when it would be better she plays risky. But the opponent stands on a risky platform. Then the voter learns about the type of one politician without incurring the costs of implementing a risky policy today. This is a welfare improvement; so we have satisfied the second claim of the proposition. The remaining claims of the proposition, that for $pr > 1$ both types of election lead to a welfare enhancement over the benchmark, follows straightforwardly from the fact that under competition the voter can screen officials due to learning.

□

Proof of Proposition 6. The proof is analogous to the proof of Proposition 3. We find the subgame perfect Nash equilibrium by backwards induction. The best response of the opposition to the first period choice of the incumbent follows: (1) when the incumbent plays risky and is successful, the opponent can only receive a positive payoff if playing risky and thereby (with positive probability) being elected; (2) when the incumbent plays risky and is unsuccessful the opponent's expected payoff from playing risky is $q_o\left[pr + (1-p)\frac{r}{2}\right] + (1-q_o)$ whilst obtaining 1 when playing safe and so the opponent plays risky if and only if only $pr > \frac{1+p}{2}$; finally (3) when the incumbent plays safe, the opponent plays risky if and only if $q_o\left[pr + (1-p)0\right) + (1-q_o)\frac{1}{2}\big] > \frac{1}{2} \Leftrightarrow pr > \frac{1}{2}$.

The optimal risk profile of the incumbent in period 1 in turn depends on the best response of the opposition. When $pr > \frac{1+p}{2}$ (opposition adopts the risky policy even when the incumbent is unsuccessful) the incumbent adopts a risky choice when:

$$p\left[r + q_o(1-p)r + q_o\frac{pr}{2}\right] + q_o(1-p)^2\frac{1}{2}) > 1 + (1-q_o)\frac{1}{2} + q_o(1-p) \tag{2}$$

Simplifying the inequality (2) we obtain that an incumbent adopts the risky policy if and only if $pr > \frac{3-q_o p^2}{4-q_o p}$ whenever $pr > \frac{1+p}{2}$.

When $pr < \frac{1+p}{2}$, so that the opponent adopts the risky policy only when then incumbent is successful, then the payoff to the incumbent from adopting the risky choice as in 2 with the only difference in the second term on the LHS of the inequality, which is now 0. The payoff from playing safe is $1 + (1-q_o)\frac{1}{2} + q_o(1-p)$ when the opponent's best response is play risky and $1 + \frac{1}{2}$ when the opponent's best response is to play safe. The latter case applies when $pr > \frac{1}{2}$ and for these

parameter values the incumbent always prefers to play safe. Instead, the former case applies when $pr < \frac{1}{2}$ and in these circumstances, the incumbent adopts a risky policy only when $pr > \frac{3+q_o-2q_op}{4-q_op}$. We conclude that the incumbent adopts a risky policy when $pr > g(p, q_o)$ where

$$g(p, q_o) = \begin{cases} pr > \frac{3-q_op^2}{4-q_op} & \text{if } pr > \frac{1+p}{2} \\ pr > \frac{3+q_o-2q_op}{4-q_op} & \text{if } pr < \frac{1+p}{2} \end{cases}$$

It is immediate that when $q_o = 0$ the previous condition reduces to the condition given in proposition 1 (non-competitive election) and when $q_o = 1$, the previous condition reduces to the condition in proposition 3 (competitive election). We note also that both branches of the equation for $g(p, q_o)$ are decreasing in q_o. It follows that, given $\frac{\partial g(p,q_o)}{\partial q_o} \leq 0$, the set of parameter values for which there are inefficiencies under elections decreases with q_o. □

Proof of Proposition 7. We claim that the voter can eradicate inefficiencies under both type of election. To prove this recall from proposition 4 that inefficiencies arise in the parameter region where $pr < 1$ and that the efficient level of first period risk involves the incumbent choosing risky if and only if $pr > \frac{1+p}{2}$. Next note that under a noncompetitive election the voter can never commit to anything other than the strategies outlined in proposition 2 unless strictly indifferent between retention or not of the incumbent. This occurs either when the incumbent plays safe or when she plays risky and is unsuccessful. Take the first case. Suppose that, in the affected parameter region where the inefficiency lies, the incumbent commits to replacing an incumbent who plays safe when $p < 1/2$ but otherwise retains the same strategy as described in proposition 2. The voter can do so since, in the absence of an active challenger, she is strictly indifferent at that point. This provides incentives for the incumbent to play risky if and only if $2pr > 1 \Leftrightarrow pr > 1/2 < \frac{1+p}{2}$ and so the inefficiency arising due to under-investment is eradicated due to the stated out of equilibrium actions of the voter. Take the second case. Suppose the voter commits to reelect an incumbent who plays safe when $p > 1/2$. Then the incumbent plays risky if $2pr > 2 \Leftrightarrow pr > 1$ so that in the region of interest the incumbent always plays safe. The inefficiency arising due to over-investment is eradicated due to the optimal equilibrium action of the voter.

In a competitive election, from proposition 4 inefficiencies arise due to the incumbent gambling on success for high levels of competence. By committing to elect the opponent when both incumbent and opponent adopt a risky policy and are successful a voter can eradicate this inefficiency. This is sequentially rational for the voter since both politicians' posterior of being competent is 1; so she is indifferent as to who implements the second period policy. Given these incentives the incumbent

chooses the risky policy if and only if $p(r + (1-p)r) > 2 - p$ which always holds. When the voter provides these incentives, the welfare attained in the area where the incumbent gambles on success in the absence of these incentives is now strictly superior to that achieved under appointment. This follows from the same arguments used above: in a competitive election the voter learns about the opponents competence without incurring the costs of policy. □

Proof of Proposition 8. Follows from the proof of proposition 7, the arguments in the main text, and the graphical depiction in figure 7. □

Proof of Proposition 9. There are four possible situations in which the incumbent and the opponent can each run on a risky or safe platform. Having the opponent stand on a risky platform and the incumbent playing safe dominates the opposite situation where the incumbent plays risky and the opponent safe or a situation in which both play safe. This is because when the opponent plays risky the voter learns about the type of one politician with no policy costs. The expected payoff to the voter in such situations is $1 + pr + (1-p)$. We compare this situation with the one in which both incumbent and opponent stand on a risky platform (even when the incumbent is unsuccessful): in this situation the voter incurs policy costs of experimentation but learns as much as possible. The expected payoff to the voter is then $2pr + (1-p)(pr + (1-p))$. From the previous payoffs we know that the voter prefers both incumbent and opponent adopting the risky alternative when $pr < h(p) = \frac{1+p-p^2}{2-p}$. From figure 8 where we plot this function we observe that when $pr > f(p)$ and $pr > h(p)$, competitive elections implement the optimal situation for the voter. Instead, when $pr \in (h(p), f(p))$ the incumbent does not adopt the risky policy though the voter would like her to do so (analogous to the fear of failure situation): the voter can provide the right incentives to achieve the desired behavior by the incumbent (reelect her with certainty when she is successful). When $pr \in (f(p), h(p))$ we are in a situation in which the incumbent adopts a risky choice when the voter would like her to not do so (analogous to the gambling on success situation): the voter can provide the right incentives to achieve the desired behavior by the incumbent (never reelect her when both the opponent and the incumbent are successful). Finally note that there is no way that a voter can induce the opponent to adopt a risky policy when $pr < 1/2$ since she cannot improve on her sequentially rational action profile: elect a successful opponent and reelect the incumbent when the opponent is not successful. □

12. Appendix B: Experimental design

We ran a total of 5 sessions with 23, 24 or 25 subjects per session. Students were recruited through the online recruitment system ORSEE (Greiner, 2003) and the experiment took place on networked personal computers in the Center for Experimental Social Sciences at Nuffield College, Oxford University in October 2011. The experiment was programmed and conducted with the software z-Tree (Fischbacher, 2007).[12]

The same procedure was used in all sessions. Instructions (see below) were read aloud and questions answered in private. Students were isolated and not allowed to communicate. The sessions consisted of 64 periods. In each period a subject was faced with a binary decision between a safe option and a risky option: the risky option's outcome depends on the outcome of a lottery. For each lottery we consider four different treatments representing a one period model (1P), a two period model (2P), a non-competitive election (NCE) and a competitive election (CE). The payoffs from choosing the safe or risky option in our experiment are simply the expected utility of our first period incumbent at adopting the safe policy or the risky policy (when anticipating a rational response by the opponent). The outcome of the lottery was not drawn until the end of the experiment when we announced the two randomly selected periods on which the subject is paid (the outcome of the lottery is then drawn and announced if the subject chose the risky option in such periods). Average payments are £12.82 (minimum payment £4, maximum payment £30.62). Session length, including waiting time and payment, is just below an hour.

INSTRUCTIONS

Thank you for agreeing to participate in our experiment. The sum of money you will earn during this experiment will be given privately to you at the end of the experiment. From now on (and until the end of the experiment) you cannot talk to any other participant. If you have a question, please raise your hand and one of the instructors will answer your questions privately. Please do not ask anything aloud!

This experiment consists of 64 periods. The rules are the same for all participants and for all periods. In each period you will have to select one of two options. Read carefully the instructions for each period as these will change from period to period.

Your choices will determine your profits in each period. At the end of the last period, the computer will randomly select 2 periods and you will earn the sum of the profits on those periods.

In Figure 9 we represent the different parameter configurations. We have 8 parameter configurations, for each parameter configuration our subjects face the four treatments in sequence (first 1P, second, 2P, third NCE and finally CE). Table 4 summarizes the exact parameter values for each of

[12]The data and programme code for the experiment are available upon request.

TABLE 4. Parameters

Combination	p	r	Treatment	Payoff(safe)	Payoff(risky|G)	Payoff(risky|R)
1	0.20	3.33	1P	7.50	25.00	0.00
1	0.20	3.33	2P	7.50	25.00	3.75
1	0.20	3.33	NCE	5.63	25.00	0.00
1	0.20	3.33	CE	6.75	23.75	1.50
2	0.20	2.30	1P	7.50	17.25	0.00
2	0.20	2.30	2P	7.50	17.25	3.75
2	0.20	2.30	NCE	5.63	17.25	0.00
2	0.20	2.30	CE	5.63	17.25	0.00
3	0.30	3.33	1P	7.50	25.00	0.00
3	0.30	3.33	2P	7.50	25.00	3.75
3	0.30	3.33	NCE	5.63	25.00	0.00
3	0.30	3.33	CE	6.38	23.13	1.31
4	0.30	2.30	1P	7.50	17.25	0.00
4	0.30	2.30	2P	7.50	17.25	3.75
4	0.30	2.30	NCE	5.63	17.25	0.00
4	0.30	2.30	CE	6.38	15.96	1.31
5	0.70	1.17	1P	7.50	8.78	0.00
5	0.70	1.17	2P	7.50	8.78	3.75
5	0.70	1.17	NCE	5.63	8.78	0.00
5	0.70	1.17	CE	4.88	7.24	0.00
6	0.70	1.03	1P	7.50	7.73	0.00
6	0.70	1.03	2P	7.50	7.73	3.75
6	0.70	1.03	NCE	5.63	7.73	0.00
6	0.70	1.03	CE	4.88	6.37	0.00
7	0.80	1.17	1P	7.50	8.78	0.00
7	0.80	1.17	2P	7.50	8.78	3.75
7	0.80	1.17	NCE	5.63	8.78	0.00
7	0.80	1.17	CE	4.50	7.02	0.38
8	0.80	1.03	1P	7.50	7.73	0.00
8	0.80	1.03	2P	7.50	7.73	3.75
8	0.80	1.03	NCE	5.63	7.73	0.00
8	0.80	1.03	CE	4.50	6.18	0.00

our combination of parameters and treatments. The first column denotes the parameter combination as depicted in figure 11; the following two columns specifically highlight the parameters r and p that correspond to each parameters configuration; the fourth column denotes the treatment (1P, 2P, NCE or CE). The fifth column denotes the payment in pounds should a subject chose the safe action (the row determines the relevant treatment and combination of parameters). Finally, the last two columns denote the payoffs from the risky option: first the payoff should a green ball be drawn, second the payoff should a red ball be drawn –there was always an urn with 10 balls and $10p$ was the number of green balls (the remaining balls were red).[13]

[13]All payoffs are multiplied by a factor of 3.75 (with respect to the values of the theoretical model) so that average payments per period are around £6.5.

In order to check whether sequencing matters, half of our subjects started the experiment facing parameter configurations with low p (combinations 1, 2, 3, and 4; in this order) whilst the other half started with combinations with high p (combinations 5, 6, 7, and 8; in this order). In the first 32 periods subjects visit the four treatments in all 8 combinations; in the last 32 periods, subjects revisit the same choices: by doing so we can also check the consistency of their actions. For instance subject 1 in session 1, needs to decide between the safe and risky option for combination 1 in the first four periods (each period corresponding to one of our four treatments), in the next four periods the subject decides among the two options when payoffs correspond to combination of parameters 2, etc... until in period 32 this subject decides among the two options in combination 4 under treatment NCE. In the following 32 periods, this subject revisits the same situations in exactly the same order.

Reassuringly we see that sequencing does not matter: regardless of subjects starting on combinations with low p (combinations 1, 2, 3, and 4) or with combinations involving high p (combinations 5, 6, 7, and 8) our three hypothesis hold true: the sign and significance of the one sided-paired difference in means (see Table 3) remain unchanged when we consider the sub-sample of subjects that started with low values of p and the sub-sample that started with high values of p.[14] In terms of consistency of our subjects' answers, 83.3% of answers in the first half of the experiment coincide with the answers given to the same situation in the second half of the experiment. We observe no pattern in the remaining observations: 10.27% modify their behavior by switching to safe and 7.43% modify their behavior by adopting risks.

[14]There is a single exception: in combination 5, the comparison between 2P and NCE is significant at 10% significance level instead of 1%

References

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

Alt, J., E. Bueno de Mesquita, and S. Rose (2011): "Accountability and Competence in Elections: Theory and Evidence from US Term Limits," *Journal of Politics*, 73(1), 1–16.

Ansolabehere, S., and J. M. Snyder (2002): "The Incumbency Advantage in U.S. Elections: An Analysis of State and Federal Offices 19422000," *Election Law Journal*, 3(1), 315–38.

Ashworth, S. (2006): "Campaign finance and voter welfare with entrenched incumbents.," *American Political Science Review*, 100, 55–68.

Ashworth, S., E. Bueno de Mesquita, and A. Friedenberg (2011): "Creating Incentives and Selecting Good Types Revisited," working paper.

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.

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.

Besley, T. J. (2006): *Principled Agents? The Political Economy of Good Government*. Oxford University Press, Oxford.

Callander, S. (2008): "A Theory of Policy Expertize," *Quarterly Journal of Political Science*, 31(2), 123–40.

——— (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.

Canes-Wrone, B., M. C. Herron, and K. W. Shotts (2001): "Leadership and pandering: a theory of executive policymaking," *American Journal of Political Science*, 45(3), 332–350.

Cao, Y., Y. Qian, and B. R. Weingast (1999): "From Federalism, Chinese Style to Privatization, Chinese Style," *Economics of Transition*, 7(1), 123–24.

Carey, J. M., R. G. Niemi, L. W. Powell, and G. F. Moncrief (2006): "The Effects of Term Limits on State Legislatures: A New Survey of the 50 States" Legislative Studies Quarterly," *Legislative Studies Quarterly*, 31(1), 105–134.

Dal Bó, E., and M. A. Rossi (2008): "Term Length and Political Performance," NBER Working Paper number 14511.

DeCicca, P., and J. D. Smith (2011): "The Long-Run Impacts of Early Childhood Education: Evidence From a Failed Policy Experiment," NBER working paper no 17085.

Erikson, R., and T. Palfrey (2006): "Equilibria in Campaign Spending Games: Theory and Data.," *American Political Science Review*, 94, 595–609.

Fearon, J. D. (1999): "Electoral Accountability and Control of Politicians: Selecting Good Types versus Sanctioning Poor Performance," in *Democracy, Accountability, and Representation*, ed. by A. Przeworski, S. C. Stokes, and B. Manin, pp. 55–97. Cambridge University Press, Cambridge.

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

Ferraz, C., and F. Finan (2009): "Electoral Accountability and Corruption: Evidence from the Audits of Local Governments," NBER Working Paper No. w14937.

Fischbacher, U. (2007): "z-Tree - Zurich Toolbox for Readymade Economic Experiments," *Experimental Economics*, pp. 171–178.

Fox, J. (2007): "Government transparency and Lawmaking," *Public Choice*, 131, 23–44.

Fox, J., and M. Stephenson (2011): "Judicial Review as a Response to Political Posturing," *American Political Science Review*, 105, 397–414.

Gordon, S. C., G. A. Huber, and D. Landa (2007): "Challenger entry and voter learning," *American Political Science Review*, 101(2), 303–320.

Gordon, S. C., and D. Landa (2009): "Do the Advantages of Incumbency Advantage Incumbents," *Journal of Politics*, 71(4), 1481–1498.

Greiner, B. (2003): "An Online Recruitment System for Economic Experiments," in *Forschung und Wissenschaftliches Rechnen*, ed. by K. Kremer, and V. Macho. Gesellschaft für Wissenschaftliche Datenverarbeitung, Göttingen: Datenverarbeitung.

Hall, P. A. (1993): "Policy Paradigms, Social Learning, and the State: The Case of Economic Policymaking in Britain," *Comparative Politics*, 25(3), 275–296.

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

Johnson, J. M., and W. M. Crain (2004): "Effects of Term Limits on Fiscal Performance: Evidence from Democratic Nations," *Public Choice*, 119(1–2), 73–90.

Kahneman, D., and A. Tversky (1979): "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, 47(2), 263–292.

Lee, D. (2008): "Randomized Experiments from Non-random Selection in U.S. House Elections," *Journal of Econometrics*, 142(2), 659–697.

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.

Maskin, E., and J. Tirole (2004): "The Politician and the Judge," *American Economic Review*, 94, 1034–1054.

Samuelson, L. (1987): "A Test of the Revealed-Preference Phenomenon in Congressional Elections," *Public Choice*, 54(2), 141–69.

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.

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.

Zaller, J. (1998): "Politicians as Prize Fighters: Electoral Selection and the Incumbency Advantage," in *Politicians and Party Politics*, ed. by J. Greer. John Hopkins University Press, Baltimore, MD.