Experimentation with Repeated Elections*

Ilwoo Hwang†

February 14, 2018

(Preliminary and Incomplete)

Abstract

When does an incumbent party have an incentive to experiment with a risky reform policy in the presence of future elections? To address this question, I study a continuous-time game between two political parties with heterogeneous preferences and a median voter. At each election, the voter chooses a party to which he gives power until the next election. Then the incumbent chooses a policy from among a safe alternative with known payoffs or two risky ones with initially unknown expected payoffs. I show that while infrequent elections are surely bad for the median voter, too frequent elections can also make him strictly worse off. When the election frequency is low, a standard agency problem arises and the incumbent party experiments with its preferred reform policy even if its outlook is not promising. On the other hand, when the election frequency is too high, in equilibrium the incumbent stops experimentation too early because the imminent election increases the incumbent's potential loss of power if it undertakes risky reform. The degree of inefficiency is large enough that too frequent elections are worse for the median voter than a dictatorship. There is an optimal frequency of elections (from the voter's perspective) that trades off the two types of inefficiencies.

Keywords: Election, strategic experimentation, political agency, learning, incumbency advantage.

JEL Classification Numbers: C73, D78, D83.

---

*I thank Hanming Fang, Camilo Garcia-Jimeno, Bård Harstad, Karam Kang, Kyungmin (Teddy) Kim, Qingmin Liu, George Mailath, Antonio Merlo, Andrew Postlewaite, Xi Weng for helpful comments and suggestions. All remaining errors are mine.

†Department of Economics, University of Miami. e-mail:ihwang@bus.miami.edu

1 Introduction

The peaceful transition of power marks a well-functioning democratic system. When an incumbent party decides whether to undertake a reform, the possibility of losing its current power affects its policy decision. In particular, if the consequences of the reform are unknown until after the reform has been implemented, then the incumbent has to consider the effects of the reform's outcomes on future elections and the action of the opposing party, which may have different preferences over policy alternatives.

For example, suppose that an incumbent party decides to implement health care reform. The diverse effects of the reform cannot be fully anticipated. Suppose also that the incumbent prefers to extend health care, while the opposing party prefers the opposite. Then the incumbent takes into account the fact that it might lose its control and the opposing party would reverse the policy. In this case, the presence of a change in power affects the incumbent's incentive to experiment with policy. Other examples of reforms whose outcome is uncertain and for which political parties have heterogeneous preferences include hawkish and dovish approaches to foreign policy, legalization of drugs, and social insurance.

In this paper, I study the incumbent party's incentives to experiment in the presence of elections, when the political parties have heterogeneous preferences over the outcomes. I address the following questions: How do repeated elections affect incentives to experiment? What is the equilibrium level of experimentation, and how does it depend on the frequency of elections? What is the socially efficient level, and is it achievable?

To address these questions, I analyze a continuous-time three-player game with two political parties and a voter. The policy experimentation process is modeled by a three-armed bandit model in which a *safe* policy yields a constant payoff, and two *risky* policies yield outcomes whose distribution, or type, is unknown. The safe policy is interpreted as a status-quo policy, while the risky policies are reform policies in different directions. At each instant, the incumbent party chooses one of the three alternatives. The parties and the voter learn the types of risky policies only through experimentation. Each risky policy is either productive or unproductive, and it is commonly known that exactly one alternative is productive while the other is unproductive. This means that one of the two mutually exclusive reform policies will turn out to be good if explored long enough. Moreover, any news shock fully reveals that the risky policy under experimentation is productive (and the other is unproductive), so all uncertainty is resolved.

Each party is biased toward a different risky alternative. This means that each reform policy has an ideological characteristic that is in accordance with one party's value. Each party gets the greatest value from its preferred risky alternative if it is productive. Furthermore, the party does not value the opposite risky alternative regardless of its types, so even if it is known that its preferred risky policy is unproductive, it prefers to choose the safe one. This payoff structure captures the loss of enthusiasm or support among the party's partisans when the opposing risky alternative is implemented.

I model elections as a Poisson arrival process. This stationarity assumption enhances the tractability of the model, while it is still enough to analyze the paper's main question: the incumbent's incentive to experiment. At each election, the voter chooses the party that will have power until the next election. The voter has unbiased preferences and prefers any productive risky alternative to the safe one. There exists a political agency problem in the sense that the voter cannot control the incumbent party while there is no election, and it is the incumbent party that chooses the policy alternatives. For the voter, the only way to control the party is to replace it with the other at the next election.

I restrict players to stationary Markov strategies with the common posterior belief as the state variable. I characterize a Markov perfect equilibrium (MPE) in which the voter elects the party whose preferred risky alternative offers a more promising belief about its productivity type. I show that there exists a unique equilibrium for a broad range of parameters. Then I analyze the efficiency of policy experimentation from the voter's perspective and conduct comparative statics with respect to the parameters such as election frequency, speed of learning and the value of a productive risky alternative.

The equilibrium shows an interesting implication in regard to the optimal frequency of elections. A common intuition is that the voter would be better off under more frequent elections, since he would have more control over the political parties. However, I show that while infrequent elections are surely bad for the voter, too frequent elections can also make him strictly worse off.

It is not surprising that if elections are infrequent, then the equilibrium exhibits inefficiency caused by political agency. If the next election is far away from now, then the incumbent party will not worry much about the future loss of power, and so its behavior is similar to that of a dictator. So the incumbent keeps experimenting with its preferred risky alternative even when the belief about its productivity type is pessimistic. The voter knows that it is better to experiment with the opposite risky policy, but he cannot control the incumbent's behavior. In this case, inefficiency decreases as election frequency increases.

A more surprising result is that there exists a different type of inefficiency under high frequency of elections. If the election frequency is greater than a certain threshold, in equilibrium the incumbent ceases to experiment at a certain point of belief. Instead, the incumbent party chooses the safe alternative and the learning stops. If elections occur too frequently, then each party would have to give up its power right after it generates the negative information. Therefore, the value of experimentation to the incumbent becomes small enough to avoid risky policy. This shows that there is another source of inefficiency from political agency: potential loss of power prevents the incumbent from conducting risky policy. I show that the degree of this inefficiency is so large that the voter is worse off under too frequent elections than under a dictatorship.

The above argument implies that there exists an optimal election frequency under which inefficiency is minimized. The frequency of elections must be high enough so that the inefficiency from a standard political agency problem is small. On the other hand, it must not be too high; otherwise it triggers a cessation of experimentation. I show that there exists a unique frequency of elections where the voter's expected payoff is maximized. Moreover, I show that the optimal frequency of elections is increasing in the value of the productive risky alternative. When the risky policy has a high value, the incumbent has enough incentive to explore the risky policy even under frequent elections.

The paper contributes to a developing literature on experimentation with multiple agents. Bolton and Harris (1999) and Keller *et al.* (2005) study a two-armed bandit problem in which different agents may choose different arms. Klein and Rady (2011) consider a similar case but assume that the expected payoffs of risky arms are negatively correlated across players, and this information structure is applied in the present paper. Bonatti and Hörner (2011) consider the case in which each agent's action is unobservable and find that the moral hazard problem leads not only to a reduction in effort but also to procrastination. In all of these papers, an informational free-riding problem leads to underinvestment in the acquisition of information. On the other hand, in the present paper the driving force for ceasing experimentation is the presence of a potential loss of control, which is crucially related to the political agency problem.[1]

The paper also contributes to the literature on political agency (Ferejohn (1986); Banks and Sundaram (1998); Besley (2004); Maskin and Tirole (2004)). These papers consider a potential moral hazard problem of elected politicians, so their results imply that it is always better for the voter to have more frequent elections. In contrast, the present paper shows that if we consider the uncertainty of policy implementation, there exists another type of political agency problem that occurs when election frequency is too high.

The paper is also related to the literature of alternating political power. Dixit and Gul (2000) use a repeated game argument to show that the presence of alternating power enables two parties to make political compromises. Aragones *et al.* (2007) develop a model of repeated elections and analyze conditions under which candidates' reputations may affect voters' beliefs over what policy will be implemented by the winning candidate of an election.

The rest of the paper is organized as follows. Section 2 describes the model. Section 3 derives the Hamilton-Jacobi-Bellman equation for a party's best response. Section 4 characterizes the Markov perfect equilibria of the non-cooperative game and derives the optimal election frequency. Section 5 argues how the voter's payoff improves in the incumbency advantage equilibrium and discusses several extensions of the model. Section 6 concludes. Some of the proofs are relegated to the Appendix.

[1] Some papers apply a model of strategic learning in the context of political economy. Strulovici (2010) considers the case in which a number of agents collectively decide which of two alternatives to choose according to some voting rule. He finds that the control-sharing effect leads to an inefficiently low level of experimentation in equilibrium. Callander and Hummel (2014) considers a two-period model to show that alternating political power can benefit voters when the policy outcome is unknown. Gul and Pesendorfer (2012) consider a model of a political campaign in which two parties of opposing interests provide costly information to voters.

2 Model

Time $t \in [0,\infty)$ is continuous. There are two political parties ($i = 1,2$) and a median voter (m). Both parties and the median voter are forward-looking and they have a common discount rate $r > 0$. At each instant, one of two political parties is determined to be an active party and chooses a policy $X_t \in \{S, R_1, R_2\}$. The other party, a passive party, cannot affect the active party's decision.

The first policy S is a *safe* policy and generates a deterministic flow payoff. There are two *risky* policies R_i ($i = 1,2$), which can be either a *productive* type or an *unproductive* type. The types of risky policies are unknown at the beginning. We will further assume a perfect negative correlation between two risky policies: it is common knowledge that exactly one risky policy is productive, while the other one is unproductive.

Payoffs for each political party are as follows. If the active party chooses to play S, then it yields a flow payoff of $s > 0$ to both parties. If the active party plays R_i and if it is unproductive, it generates zero payoff. If R_i is productive, then it pays a lump-sum payoff h at random times *only to party* i. Party j gets zero payoff from R_i regardless of its type. These heterogeneous preferences can be interpreted such that the party j is biased toward the risky action R_j, so that party j does not value the outcomes from R_i. Assume that the lump-sum arrival times correspond to jumping times of a Poisson process with intensity $\lambda > 0$. Then $g = h\lambda$ is the expected payoff to party i per unit of time, conditional on R_i being productive. We assume $g > s > 0$, so party i strictly prefers R_i to S and S to R_j if R_i is productive, and strictly prefers S, if it is unproductive, to R_i and R_j.

While each political party is biased toward the outcomes of one risky policy, the median voter is unbiased toward both risky policies and hence prefers *any* productive risky policy. That is, he gets the expected payoff of g from any productive risky policy and the flow payoff of s from a safe policy.

I model elections as Poisson arrivals. At random time, which corresponds to jumping times of a Poisson process with arrival rate $\xi > 0$, the median voter chooses one of the two parties to be the active party. Once a party is chosen, then it is guaranteed to have control over the action choices until the next election. The election process and the lump-sum payoff process of a productive risky policy are independent. The types of the risky policies stay the same at every regime change, that is, Nature conducts a random draw only once at the beginning of the game. Finally, we assume no private information: both parties can observe the active party's choice of action and the resulting outcome.

Let $\{\sigma_{i,t}\}_{t\geq 0}(i=1,2)$ and $\{\sigma_{M,t}\}_{t\geq 0}$ be the actions of the parties and the median voter, where $\sigma_{i,t} \in \{S, R_1, R_2\}$ and $\sigma_{M.t} \in [0,1]$ (probability of choosing party 1) is measurable with respect to the information available at time t. Let $\{\iota_t\}_{t\geq 0}$ $(\iota_t \in \{1,2\})$ be a stochastic process of the active party, which is determined by Poisson arrivals of the elections and $\{\sigma_{M,t}\}_{t\geq 0}$. Then a policy decision rule is a stochastic process $\{X_t\}_{t\geq 0}$, where $X_t = \sigma_{\iota_t,t}$ is an action taken by the active party at time t.

Let p_t be the common posterior belief at time t that R_1 is productive. Let $\{\mathscr{F}_t\}_{t\geq 0}$ be a filtration generated by $\{X_t\}_{t\geq 0}$ and the corresponding outcome process, then the stochastic process $\{p_t\}_{t\geq 0}$ is adapted to $\{\mathscr{F}_t\}_{t\geq 0}$, and p_t evolves according to Bayes' rule. The posterior belief jumps up to one once there is a breakthrough on R_1, and jumps down to zero if a breakthrough on R_2 is observed. In either case, learning is complete and p_t stays the same. If there has been no breakthrough until t, then p_t obeys the following differential equation:

$$\dot{p}_t = p_t(1-p_t)\lambda(\mathbf{1}_{X_t=R_2} - \mathbf{1}_{X_t=R_1}).$$

Note that $\dot{p}_t < 0 (> 0)$ when the active party plays $R_1(R_2)$ and no breakthrough is discovered.

Party 1's total discounted expected payoff, expressed in per-period units, can be written as

$$\mathrm{E}_0\left[\int_0^\infty re^{-rt}[\mathbf{1}_{X_t=R_1}\cdot p_t g + \mathbf{1}_{X_t=S}\cdot s]dt\right]$$

where the expectation is taken over the stochastic processes $\{X_t\}_{t\geq 0}$ and $\{p_t\}_{t\geq 0}$. Similarly, the payoff for party 2 is

$$\mathrm{E}_0\left[\int_0^\infty re^{-rt}[\mathbf{1}_{X_t=R_2}\cdot (1-p_t)g + \mathbf{1}_{X_t=S}\cdot s]dt\right],$$

and

$$\mathrm{E}_0\left[\int_0^\infty re^{-rt}[\mathbf{1}_{X_t=R_1}\cdot p_t g+\mathbf{1}_{X_t=R_2}\cdot(1-p_t)g+\mathbf{1}_{X_t=S}\cdot s]dt\right]$$

for the median voter.

A sequential equilibrium is called a Markov perfect equilibrium (MPE) if the agents of the game play stationary Markov strategies with the common posterior belief $p_t \in [0,1]$ as a state variable. Party i's strategy is then given by a function $\sigma_i : [0,1] \to \{S, R_1, R_2\}$, and the voter's strategy is given by $\sigma_M : [0,1] \to [0,1]$.

Let Σ^* be the set of strategy profiles where the median voter chooses the candidate whose preferred reform policy has a more optimistic belief. In other words, the median voter's strategy is $\sigma_M(p) = 1$ whenever $p > 1/2$ and $\sigma_M(p) = 0$ whenever $p < 1/2$. In this paper, I focus on Markov perfect equilibria in Σ^*.

Observe that in any equilibrium in Σ^*, each party never plays the reform policy that the opposite party prefers, that is, party i always chooses either S or R_i. This is because for party i, playing R_j gives the same amount of information as R_i while generating a zero flow payoff. In the rest of the paper, I will define $k_i(p) \in \{0,1\}$ as the probability that party i chooses its preferred reform policy.

3 Preliminary Analysis

3.1 Single party's problem

Before I analyze the model, consider a benchmark case where there is no election, that is, $\xi = 0$. In this case, only one party is active for all $t \in [0,\infty)$ and it faces the single decision-maker problem described in Keller *et al.* (2005). Suppose party 1 is always active, and let $V_1(p)$ be its value as a function of posterior belief p. Then it solves the following Hamilton-Jacobi-Bellman equation:

$$V_1(p) = s + \max_{k_1\in\{0,1\}} k_1\left\{-s+pg+\frac{p\lambda}{r}\{g-V_1(p)-(1-p)V_1'(p)\}\right\}. \tag{1}$$

The first part of the maximand corresponds to action S, and the second corresponds to R_1. The effect of R_1 on the value of party 1 can be decomposed into three elements: (i) an expected flow payoff pg, (ii) a jump in value function when party 1 discovers a breakthrough on R_1, captured by $\frac{p\lambda}{r}(g - V_1(p))$, and (iii) a decrease in value function when no breakthrough is observed, captured by $-\frac{\lambda}{r}p(1-p)V_1'(p)$. The difference between expected flow payoffs from S and R_1, which is

$$c_1(p) \equiv s - pg,$$

is called the *opportunity cost* of experimentation for party 1. The sum of the second and third elements,

$$b_1(p) \equiv \frac{p\lambda}{r}\{g - V_1(p) - (1-p)V_1'(p)\},$$

is called the *value of experimentation* of party 1. Then party 1 experiments if and only if $b_1(p) > c_1(p)$. The value of information $b_1(p)$ is nonnegative in the single party problem. However, I show that if $\xi > 0$, it can be negative for a range of p under some equilibria.

If party 1 were myopic, i.e., merely maximizing current flow payoffs, then it plays the risky action if and only if $c_1(p)$ is negative. So party 1 plays the threshold strategy with threshold $p^m = \frac{s}{g}$. If it were forward-looking, then it values the information from a risky action to use it for future decisions. In this case, the optimal decision rule is to play a risky action if and only if $b_1(p) > c_1(p)$, so it uses the optimal single party threshold

$$p^0 = \frac{\mu s}{\mu g + (g-s)} < p^m, \tag{2}$$

where $\mu = \frac{r}{\lambda}$.

3.2 Hamilton-Jacobi-Bellman equation

Now I introduce elections by assuming $\xi > 0$. Let $V_1(p)(W_1(p))$ be the value function of party 1 when it is active (passive). Then for any open interval of beliefs where the actions of party 2 and the median voter are constant, party 1's payoff function is differentiable and solves the following set of Hamilton-Jacobi-Bellman (HJB) equations:[2]

$$
\begin{aligned}
V_1(p) &= s-(1-\sigma_M(p))\frac{\xi}{r}(W_1(p)-V_1(p)) \\
&\quad + \max_{k_1\in\{0,1\}} k_1\left\{-s+pg+\frac{p\lambda}{r}\{g-V_1(p)-(1-p)V_1'(p)\}\right\},
\end{aligned} \tag{3}
$$

$$
W_1(p) = \begin{cases} \sigma_M(p)\frac{\xi}{r}(V_1(p)-W_1(p))+s, & \text{if } k_2(p)=0,\\ \sigma_M(p)\frac{\xi}{r}(V_1(p)-W_1(p))+\frac{(1-p)\lambda}{r}\{W_1(0)-W_1(p)+pW_1'(p)\}, & \text{if } k_2(p)=1. \end{cases} \tag{4}
$$

Similar to the single decision-maker problem, party 1, when active, faces the trade-off between the opportunity cost and the value of experimentation. However, there exists an election after which party 1 can lose its power and become a passive player if the median voter chooses party 2. The first term of the equation for $V_1(p)$ represents such a possible regime change. Note that party 1's problem is essentially the same as that of the single decision-maker for the range of beliefs where the median voter chooses party 1. Since it cannot affect the action choice when it is passive, there is no maximization problem in the formula of $W_1(p)$, and $W_1(p)$ depends on the opponents' strategy. The first term of $W_1(p)$ disappears when the median voter chooses party 2.

For the median voter, let $Z_m(p,i)$ be the value function of the median voter when there is no election and when party i is active, and $V_m(p)$ be the value function at the time of election. Then the HJB equations

[2] Similarly, for any open interval of beliefs where party 1's action is constant, $V_2(p)$ and $W_2(p)$ are differentiable and they solve

$$
\begin{aligned}
V_2(p) &= s-\sigma_M(p)\frac{\xi}{r}(W_2(p)-V_2(p)) \\
&\quad + \max_{k_2\in\{0,1\}} k_2\left\{-s+(1-p)g+\frac{(1-p)\lambda}{r}\{V_2(0)-V_2(p)+pV_2'(p)\}\right\},
\end{aligned}
$$

$$
W_2(p) = \begin{cases} (1-\sigma_M(p))\frac{\xi}{r}(V_2(p)-W_2(p))+s, & \text{if } k_1(p)=0,\\ (1-\sigma_M(p))\frac{\xi}{r}(V_2(p)-W_2(p))+\frac{p\lambda}{r}\{W_2(1)-W_2(p)-(1-p)W_2'(p)\}, & \text{if } k_1(p)=1. \end{cases}
$$

are given by

$$
\begin{aligned}
Z_M(p,i) \quad = \quad & \frac{\xi}{r}(V_M(p) - Z_M(p,i)) \\
& + \mathbf{1}_{\{\sigma_i(p)=S\}} \cdot s \\
& + \mathbf{1}_{\{\sigma_i(p)=R_1\}} \cdot \{pg + \frac{p\lambda}{r}\left\{Z_M(1,i) - Z_M(p,i) - (1-p)Z'_M(p,i)\right\}\} \\
& + \mathbf{1}_{\{\sigma_i(p)=R_2\}} \cdot \{(1-p)g + \frac{(1-p)\lambda}{r}\left\{Z_M(0,i) - Z_M(p,i) + pZ'_M(p,i)\right\}\},
\end{aligned} \tag{5}
$$

$$
V_M(p) \quad = \quad Z_M(p,2) + \max_{\sigma_M \in [0,1]} \sigma_M [Z_M(p,1) - Z_M(p,2)]. \tag{6}
$$

Note that there is no maximization problem in the expression for $Z_M(p,i)$. This is because the median voter does not take any action between elections, so only the active party's policy choice affects the value of $Z_M(p,i)$. When the election comes, the median voter chooses the party that gives him higher $Z_M(p,i)$.

When the uncertainty is resolved, i.e., when $p = 0$ or 1, there exists a dominant strategy for each player. Since there is no uncertainty at those beliefs, both parties choose the myopically optimal action. Therefore, $(\sigma_1(0), \sigma_2(0)) = (S, R_2)$, and $(\sigma_1(1), \sigma_2(1)) = (R_1, S)$. Moreover, the median voter chooses the party that is biased toward the productive risky arm, so $\sigma_M(0) = 0$ and $\sigma_M(1) = 1$. Using this, the values of V_1 and W_1 in the certainty case are calculated and given by

$$
\begin{aligned}
V_1(1) = g, \qquad & V_1(0) = (1-\chi)s \\
W_1(1) = \chi g + (1-\chi)s, \qquad & W_1(0) = 0.
\end{aligned}
$$

where $\chi = \frac{\xi}{\xi + r}$.

3.3 Preliminary Observations

First, we prove lemmas which makes the analysis simpler. The first lemma states that in the belief range where there is no chance of losing power, each party plays as if it were a single decision-maker.

Lemma 1. *In any MPE in Σ^*, the action of party 1 (resp. party 2) at $p > 1/2$ (resp. $p < 1/2$) is the same as that of the optimal decision rule of the single decision-maker problem.*

Proof. It is sufficient to consider party 1. Since $\sigma_M^*(p) = 1$ for $p > 1/2$, party 1's HJB equations in this range of beliefs given σ_M^* are equal to those of the single decision-maker problem (1). Therefore, the corresponding value functions and the boundary conditions are also the same, which leads to the same optimal decision rule in the equilibrium. □

The next lemma shows that each party has the dominant strategy when the belief is close to certainty.

Lemma 2. *In any MPE in Σ^*, 1) party 1 (resp. party 2) chooses R_1 (resp. R_2) at a belief close to one (resp. zero), and 2) chooses S at a belief close to zero (resp. one).*

Proof. The first part is straightforward from Lemma 1. Suppose to the contrary that there exists an MPE where party 1 chooses R_1 for any $p > 0$. Then by the first part of the proof, there exists $\underline{p} \in (0, 1/2)$ such that $k_2(p) = 1$. Since $\sigma_M(p) = 0$ for $p < \underline{p}$, we can derive the functional form of value function $V_1(p)$ for $p \in (0, \underline{p})$. Then a simple calculation shows that $\lim_{p \to 0} V(p) = 0$ (it is sufficient to show that $C_6 = 0$), but then party 1 can deviate to S to get the payoff of $(1 - \chi)s$. Contradiction. □

Lemma 2 implies that in any MPE, there exists at least one threshold belief point $p_1 \in (0, 1)$ where party 1 switches action from R_1 to S, that is, there exists $\varepsilon > 0$ such that $k_1(p) = 1$ if $p \in (p_1, p_1 + \varepsilon)$ and $k_1(p) = 0$ if $p = p_1$. Define p_1^* be the greatest threshold point of party 1, and let p_2^* be the smallest threshold point of party 2.

4 Equilibrium

In this section, I fully characterize the Markov perfect equilibria in Σ^*. There are three classes of MPEs for which the equilibrium outcome is qualitatively different. I show that each class of MPE appears in a different range of parameters.

4.1 Low stake

The following theorem (whose proof is in the appendix) states that in the case of a small stake (low value of g/s), the there exists a unique MPE which generates an efficient outcome from the voter's perspective. Recall that p^0 is an optimal threshold of the single decision-maker problem.

Proposition 1. *If $\frac{g}{s} < \alpha_0 \equiv \frac{1+2\mu}{1+\mu}$, then a strategy profile in Σ^* is an MPE if and only if $\sigma_1^{-1}(R_1) = (p^0, 1]$* and $\sigma_2^{-1}(R_2) = [0, 1 - p^0)$.

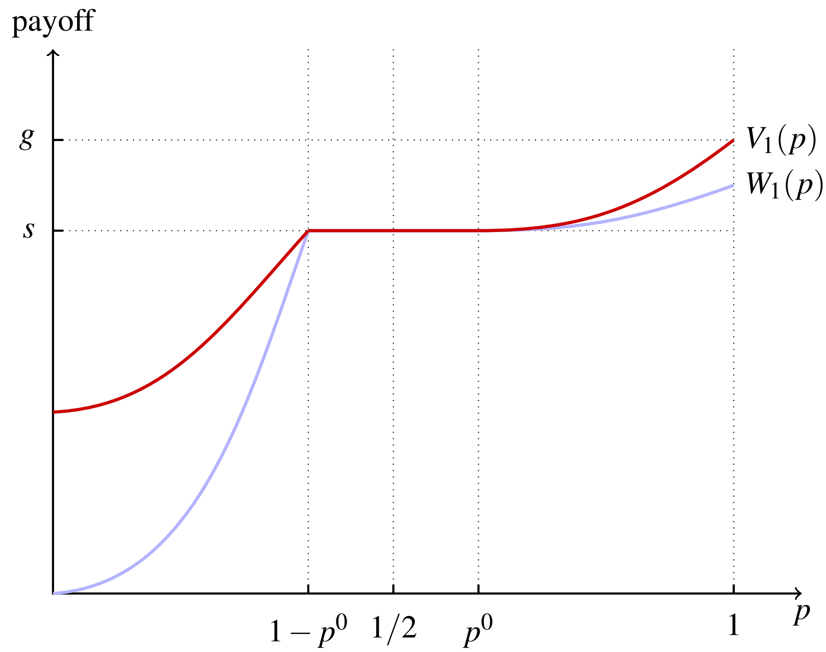

Figure 1: Party 1's payoff function in the MPE of the low-stake case (parameter values: $g/s = 1.2, r = 0.05, \lambda = 0.1, \xi = 0.05$).

From (2) it is easy to see that $p^0 > \frac{1}{2}$ if and only if $\frac{g}{s} < \alpha_0$. Therefore, in the above equilibrium, once the belief falls in the range $[1-p^0, p^0]$, both parties choose the safe policy and the reform policy is never explored.

In this equilibrium, each party's equilibrium strategy is the same as the that of the single decision-maker's problem in Section 3.1. This is because given the median voter's strategy, there is essentially no strategic interaction between the two parties. Note that the incumbent is in the "safe region" if its preferred reform policy has a more optimistic belief, because the voter would reelect the incumbent if there was an election at that instant. Then the incumbent in its safe region chooses the policy as if it is a single decision-maker. Here, each party stops experimentation in the safe region; hence, its optimal strategy does not depend on the opponent party's strategy. The similar intuition explains the fact that the upper bound α_0 on the size of the stake does not depend on the election frequency (ξ).

Figure 1 describes party 1's payoff functions $V_1(p)$ (deak red line) and $W_1(p)$ (bright blue line) in the equilibrium in the low-stake case. Note that $W_1(p)$ is less than or equal to $V_1(p)$ for any belief point, and the difference between the two functions captures the cost of losing control. Both $V_1(p)$ and $W_1(p)$ are equal to s in the middle range of the belief space, since both parties choose the safe policy.

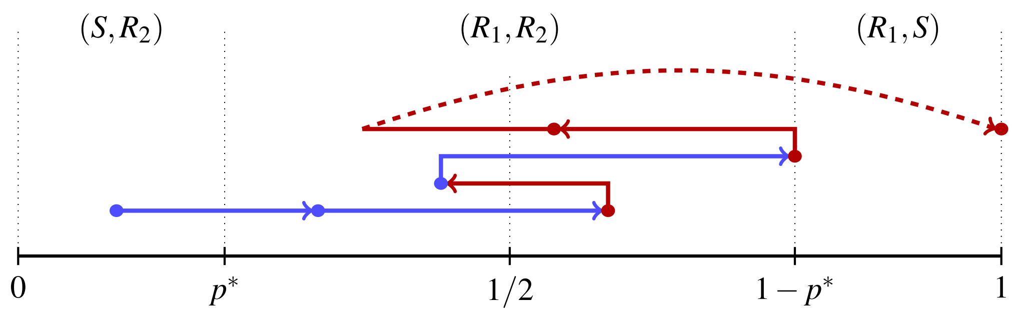

Figure 2: An equilibrium outcome path in the infrequent election case

In the low-stake case, the unique equilibrium outcome is optimal for the median voter, and there is no inefficiency from the political agency problem.

4.2 High stake case: Infrequent elections

Now consider the high-stake case where $\frac{g}{s} \geq \alpha_0$. In this case, the optimal threshold p^0 of the single decision-maker problem becomes strictly less than $1/2$. So playing p^0 induces a nontrivial strategic interaction between the parties, so the equilibrium strategy profile would differ from that of the single decision-maker's problem. There are two types of equilibria in the high-stake case: one emerges in the case where the election frequency is low and the other emerges in the frequent elections case. Both types of equilibria show inefficient outcomes (from the voter's perspective), but the underlying forces for the inefficiency is different in each type of equilibrium.

Proposition 2. *There exists $\alpha_3(\chi)$ such that if $\frac{g}{s} > \alpha_3$, there exists a unique $p^* < \frac{1}{2}$ such that a strategy profile in Σ^* is an MPE if and only if* $\sigma_1^{-1}(R_1) = (p^*, 1]$ and $\sigma_2^{-1}(R_2) = [0, 1-p^*)$. *Moreover, $\alpha_3(\chi)$ is increasing in χ.*

In this equilibrium, each party chooses the reform policy even when the belief is unfavorable to its preferred reform policy. Hence failure to find a breakthrough eventually leads to political turnover. Since the voter always elects the party that would conduct one of the reform policies, (with probability one) there is a breakthrough in the reform policy in finite time; hence, the uncertainty is resolved.

Figure 2 describes an equilibrium outcome path of the infrequent election case in terms of belief dynamics. The dark red line (bright blue line) represents the belief dynamic when party 1 (party 2) is a ruling party, and a circle represents an election. In this case, the prior belief p_0 is less than a half, so

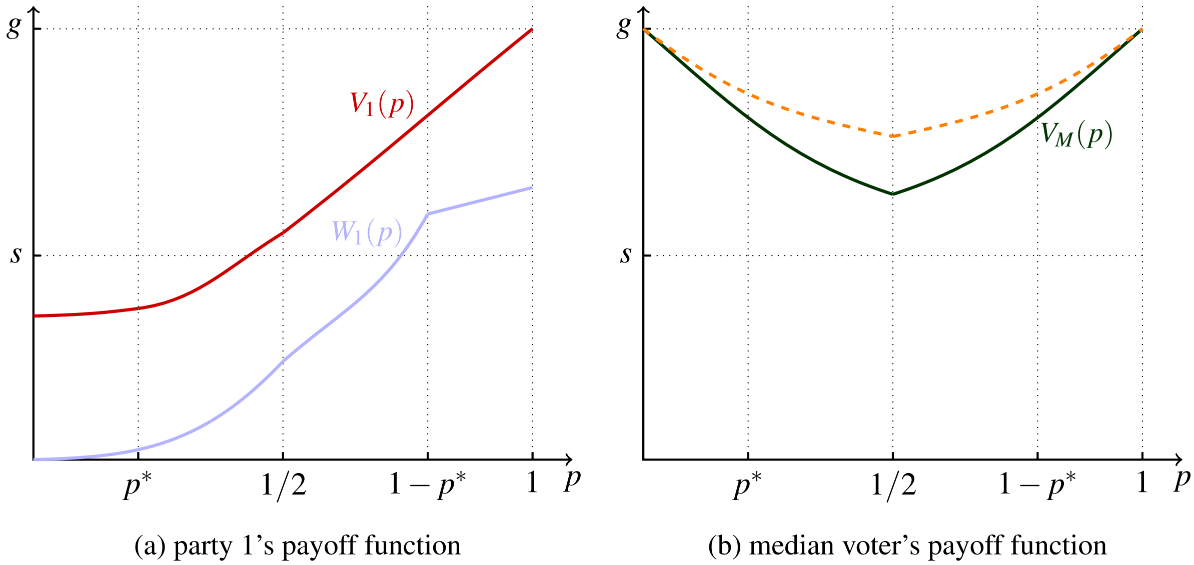

Figure 3: Equilibrium behavior of the MPE in the infrequent election case (parameter values: $g/s = 2.2, r = 0.05, \lambda = 0.1, \xi = 0.02$).

initially the median voter elects party 2, which chooses its preferred reform policy (R_2). Once the belief goes above a half, the voter chooses party 1 at the election. There will be successive political turnovers until one of the parties receives a breakthrough from its preferred reform policy.

Figure 3 describes party 1's payoff functions (upper panel) and the median voter's payoff function (lower panel) in the equilibrium in Theorem 2. In the lower panel, the green line represents the median voter's expected payoff function $V_M(p)$. The dashed yellow line is the median voter's payoff under no political agency problem, that is, the payoff when he could choose the policy by himself. It turns out that for any $p \in (0,1)$, the voter's expected payoff is strictly less than the payoff with no agency problem. In the equilibrium of the infrequent election case, the inefficiency comes from the suboptimal choice of a risky policy by the ruling party.

4.3 High stake case: Frequent elections

If the election is frequent, there exists another type of inefficiency. In the equilibrium of the frequent election case, the incumbent stops experimentation too early because the imminent election increases the incumbent's potential loss of power if it undertakes risky reform. The degree of inefficiency is large enough that too frequent elections are worse for the median voter than a dictatorship.

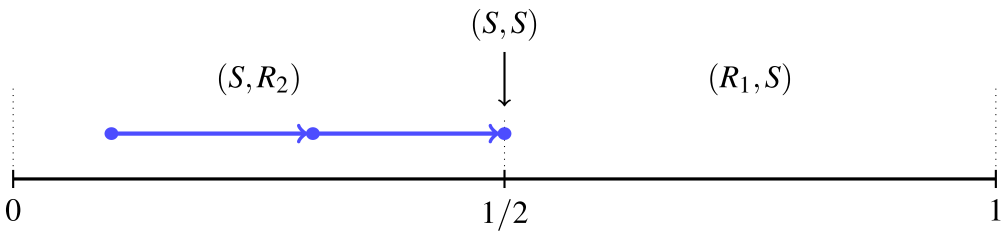

Figure 4: An equilibrium outcome path in the frequent election case

In the frequent election case, there exist three types of equilibria that appear in the different range of parameters. However, all types of equilibria share the common feature that the expected length of experimentation is shorter than the efficient level, and uncertainty is not resolved with positive probability.

Proposition 3. *There exists $\alpha_1(\chi)$ such that for $\frac{g}{s} \in [\alpha_{0,}\alpha_1]$, a strategy profile in Σ^* is an MPE if and only if $\sigma_1^{-1}(R_1) = (\frac{1}{2}, 1]$ and $\sigma_2^{-1}(R_2) = [0, \frac{1}{2})$. Furthermore, $\alpha_1(\chi)$ is increasing in χ.*

Figure 4 describes a possible belief path in the equilibrium of Theorem 3. In this equilibrium, each incumbent initially chooses its preferred reform policy. But if a breakthrough has not been discovered until the belief reaches 1/2, then the incumbent stops experimentation and switches to the safe policy. Since the opponent party would also choose the safe policy at $p = 1/2$, replacing the incumbent does not help in terms of more reform policy. This implies that for any $p \in (0,1)$, the uncertainty is not resolved with positive probability.

Figure 5 describes party 1's payoff function (upper panel) and the median voter's payoff function (lower panel) in the equilibrium in Theorem 3. Note that in the upper panel, $V_1(p)$ and $W_1(p)$ have the same value at $p = 1/2$ as both parties play a safe policy at that point. In the lower panel, the median voter's expected payoff function $V_M(p)$ (green line) hits the value s at the belief $1/2$ as there is no experimentation at that point. Similar to Figure 3, the degree of inefficiency is captured as the distance between the dashed yellow line and the green line. In the equilibrium of the frequent election case, the inefficiency comes from underinvestment in the reform policy. Note that all parameter values used for Figures 3 and 5 are the same except the election frequency, and that the voter's expected payoff function is lower in the frequent election case (Figure 5). In the next subsection, I will discuss more about the relationship with the election frequency and the voter's welfare.

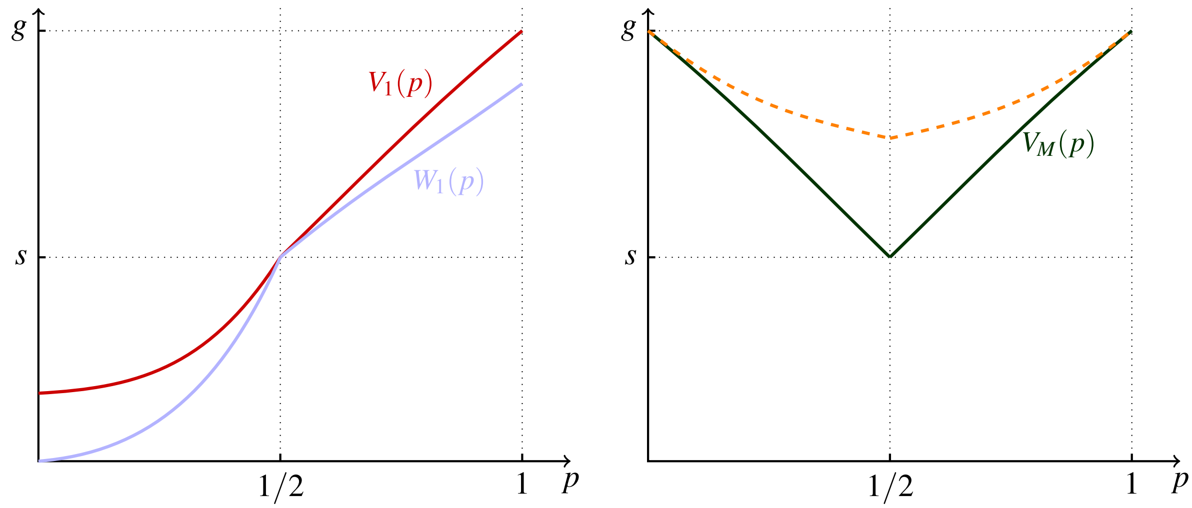

(a) party 1's payoff function

(b) median voter's payoff function

Figure 5: Equilibrium behavior of the MPE in the frequent election case (parameter values: $g/s = 2.2, r = 0.05, \lambda = 0.1, \xi = 0.2$).

4.4 Voter's welfare and optimal election frequency

The equilibrium analysis in the previous subsection suggests that there is an optimal frequency of elections (from the voter's perspective) that trades off the two types of inefficiencies.

Proposition 4. *Suppose $g/s > \alpha_0$. Then for any $p \in (0,1)$,*

1. *if the parameter values are in the 'infrequent elections' range, $V_M(p;\xi)$ is increasing in the election frequency (ξ);*
2. *$V_M(p)$ in the MPE of the frequent election case is smaller than the one in the MPE of the infrequent election case;*
3. *$V_M(p)$ in the MPE of the frequent election case is smaller than the one in the dictatorship case ($\xi = 0$).*

If the parameter values are in the infrequent elections range, the degree of inefficiency decreases as election frequency increases. However, too frequent elections would result in the worst outcome from the perspective of the median voter, as the parties stop experimentation in the equilibrium of the frequent election case. In fact, the median voter's expected payoff in the frequent election case is lower than the one where there is no election, as the length of experimentation is shorter.

5 Incumbency advantage

For the frequent election case, I conjecture that that efficiency can be restored by giving an advantage to the incumbent in the election.

Proposition 5. *Suppose* $\frac{g}{s} \in [\alpha_{0,}\alpha_1]$. *Then there exists* $x_1, x_2 > 0$ *such that the following strategy profile is an MPE:*

1. On the equilibrium path, party i plays threshold strategy with threshold $\frac{1}{2} + x_1 \cdot (-1)^i$, *and the median voter plays*

$$\sigma_M(p,i) = \begin{cases} 1 & \text{if } i = 1, p \geq \frac{1}{2} - x_1 \text{or } i = 2, p > \frac{1}{2} + x_1 \\ 0 & \text{otherwise.} \end{cases}$$

2. If the median voter deviates, then the agents play a Markovian profile with $\sigma_1^{-1}(R_1) = (\frac{1}{2}, 1]$ and $\sigma_2^{-1}(R_2) = [0, \frac{1}{2})$.

3. If the party i deviates, then the agents play agents play a Markovian profile with $\sigma_1^{-1}(R_1) = (\frac{1}{2} + x_2 \cdot (-1)^{i-1}, 1]$, $\sigma_2^{-1}(R_2) = [0, \frac{1}{2} + x_2 \cdot (-1)^{i-1})$, and $\sigma_M^{-1}(1) = (\frac{1}{2} + x_2 \cdot (-1)^{i-1}, 1]$.

The above equilibrium is non-Markovian where the voter chooses the incumbent at the neighbor of $p = 1/2$. By giving advantage to the incumbent, the voter can induce each party to experiment and induce endogenous political turnover. I conjecture that the above profile is still an equilibrium even if the election frequency is arbitrarily high, so the median voter can approximately achieve first-best.

In the incumbency advantage equilibrium, the voter is more generous to the incumbent in the sense that he may reelect the incumbent even when its preferred risky alternative is less promising than the opposite one. Knowing that, the incumbent experiments aggressively with its preferred risky policy even under frequent elections. Therefore, the incumbency advantage strategy can introduce frequent switches of power without causing the cessation of the experimentation with the risky alternatives, which is optimal from the voter's perspective. This result provides a normative argument for the incumbency advantage and contributes to previous positive arguments about the incumbency advantage.[3]

[3]For the positive argument of incumbency advantage, see Samuelson (1987) and Ashworth and de Mesquita (2008).

6 Conclusion

In this paper, I study a continuous-time game between two political parties with heterogeneous preferences and a median voter. At each election, the voter chooses a party to which he gives power until the next election. Then the incumbent chooses a policy from among a safe alternative with known payoffs or two risky ones with initially unknown expected payoffs. I show that while infrequent elections are surely bad for the median voter, too frequent elections can also make him strictly worse off. When the election frequency is low, a standard agency problem arises and the incumbent party experiments with its preferred reform policy even if its outlook is not promising. On the other hand, when the election frequency is too high, in equilibrium the incumbent stops experimentation too early because the imminent election increases the incumbent's potential loss of power if it undertakes risky reform. The degree of inefficiency is large enough that too frequent elections are worse for the median voter than a dictatorship. There is an optimal frequency of elections (from the voter's perspective) that trades off the two types of inefficiencies.

The equilibrium outcome in the infrequent election case in Figure 2 suggests that there is a correlation between the time in power and the electability of the incumbent. More formally, let T_0 be the length of time in which the current ruling party has been in power, and let T_{-k} be the length of time in which the kth previous ruling party had been in power. Furthermore, let $\tilde{\pi}$ be the probability that the current ruling party will be reelected. Then I conjecture that under the equilibrium in the infrequent election case, if there has been no breakthrough, $\tilde{\pi}$ is decreasing (increasing) in T_{-k} for k even (odd).

Appendices

A Explicit Solutions to HJB Equations

Given a Markov strategy of party 2 and the median voter (σ_2, σ_M), standard arguments imply that on any open interval where k_2 and σ_M are constant, party 1's payoff functions (V_1, W_1) from playing a best response are continuously differentiable and solve the system of HJB equations

$$V_1(p) = s - (1-\sigma_M(p)) \cdot \tau_1(p; V_1, W_1) + \max_{k_1 \in [0,1]} k_1 (b_1(p; V_1) - c_1(p)), \tag{7}$$

$$W_1(p) = s + \sigma_M(p) \cdot \tau_1(p; V_1, W_1) + k_2(p)(\beta_1(p; W_1) - s), \tag{8}$$

where

- $\tau_1(p; V_1, W_1) = \frac{r}{\xi}(V_1(p) - W_1(p))$: value change from loss of control;
- $b_1(p; V_1) = \frac{p}{\mu}\{g - V_1(p) - (1-p)V_1'(p)\}$: value of party 1's experimentation;
- $\beta_1(p; W_1) = \frac{(1-p)}{\mu}\{-W_1(p) + pW_1'(p)\}$: value of opponent's experimentation;
- $c_1(p) = s - pg$: opportunity cost of experimentation,

and $\mu = r/\lambda$ is the inverse of the effective success rate of the reform policy. Party 2's value functions $V_2(p)$ and $W_2(p)$ satisfy a similar HJB equation, with $1-p$ replacing p.

For a range of beliefs where the Markov profile (k_1, k_2, σ_M) is constant, we use (7) and (8) to solve the explicit solutions for $V_1(p)$ and $W_1(p)$. First, consider the case where $\sigma_M(p) = 1$. Then (7) is written as $V_1(p) = s + \max_{k_1 \in [0,1]} k_1 (b_1(p; V_1) - c_1(p))$, and solving the differential equation for each value of k_1 gives

$$V_1(p) = \begin{cases} s & \text{if } k_1 = 0, \\ gp + C_2 f(p) & \text{if } k_1 = 1, \end{cases}$$

where C_2 is an integration constant and $f(p) = (1-p)^{\mu+1} p^{-\mu}$. Then (8) implies that the solution is.

$$W_1(p) = \begin{cases} s + \tau_1(p; V_1, W_1) & \text{if } k_2 = 0, \\ s + \tau_1(p; V_1, W_1) + k_2(p)(\beta_1(p; W_1) - s) & \text{if } k_2 = 1, \end{cases}$$

(k_1,k_2)	$V_1(p)$	$W_1(p)$
$(0,0)$	s	s
$(0,1)$	s	$\frac{\hat{\mu}\chi s}{\hat{\mu}+1}+\frac{\chi s}{\hat{\mu}+1}p+C_1\hat{f}(1-p)$
$(1,0)$	$gp+C_2f(p)$	$(1-\chi)s+\chi(gp+C_2f(p))$
$(1,1)$	$gp+C_2f(p)$	$\chi gp+\frac{\hat{\mu}\chi}{\hat{\mu}+\mu+1}C_2f(p)+C_3\hat{f}(1-p)$

(a) Value functions when $\sigma_M(p)=1$

(k_1,k_2)	$V_1(p)$	$W_1(p)$
$(0,0)$	s	s
$(0,1)$	$(1-\chi)s+\chi C_5f(1-p)$	$C_5f(1-p)$
$(1,0)$	$\chi s+\frac{g-\chi s}{\hat{\mu}+1}p+C_4\hat{f}(p)$	s
$(1,1)$	$\frac{\mu+1}{\hat{\mu}+1}gp+\frac{\hat{\mu}\chi}{\hat{\mu}+\mu+1}C_5f(1-p)+C_6\hat{f}(p)$	$C_5f(1-p)$

(b) Value functions when $\sigma_M(p)=0$

Table 1: Explicit solutions to HJB equations for party 1

Solving the differential equations for each pair of (k_1,k_2), we have

$$W_1(p)=\begin{cases} s & \text{if } (k_1,k_2)=(0,0),\\ \frac{\hat{\mu}\chi s}{\hat{\mu}+1}+\frac{\chi s}{\hat{\mu}+1}p+C_1\hat{f}(1-p) & \text{if } (k_1,k_2)=(0,1),\\ (1-\chi)s+\chi(gp+C_2f(p)) & \text{if } (k_1,k_2)=(1,0),\\ \chi gp+\frac{\hat{\mu}\chi}{\hat{\mu}+\mu+1}C_2f(p)+C_3\hat{f}(1-p) & \text{if } (k_1,k_2)=(1,1),\end{cases}$$

where C_1 and C_3 are integration constants, $\hat{\mu}=\mu/(1-\chi)$, and $\hat{f}(p)=(1-p)^{\hat{\mu}+1}p^{-\hat{\mu}}$. Table 1a shows the explicit solutions when $\sigma_M(p)=1$.

For the range of beliefs in which $\sigma_M(p)=0$, we first solve (8) to obtain explicit solution of $W_1(p)$ for each value of $k_2=0,1$. Then we plug the solution into equation (7) to get the solution of $V_1(p)$ for each value of $k_1=0,1$. Table 1b shows the explicit solutions when $\sigma_M(p)=0$, with integration constants C_4,C_5, and C_6.

As we shall see, the integration constants $C_k(k=1,\dots,6)$ are determined by boundary conditions, such as value-matching conditions and smooth-pasting conditions. With some abuse of notation, we use

(k_1,k_2)	$Z_M(p,1)$	$Z_M(p,2)$
$(0,0)$	s	s
$(0,1)$	s	$\frac{p+\hat{\mu}}{\hat{\mu}+1}\chi s+\frac{\mu+1}{\hat{\mu}+1}g(1-p)+D_2\hat{f}(1-p)$
$(1,0)$	$gp+D_1f(p)$	$(1-\chi)s+\chi(gp+D_1f(p))$
$(1,1)$	$gp+D_1f(p)$	$\chi gp+\frac{\mu+1}{\hat{\mu}+1}g(1-p)+$ $\chi D_1\frac{\hat{\mu}}{\hat{\mu}+\mu+1}f(p)+D_3\hat{f}(1-p)$

Table 2: Explicit solutions to HJB equations for the median voter, when $\sigma_M(p)=1$

the same notation C_k for different types of equilibria, although the value of C_k varies in each equilibrium.

A similar calculation shows that given a Markov strategy $(\sigma_1,\sigma_2,\sigma_M)$, on any open interval where values of k_1, k_2 and σ_M are constant, the median voter's payoff functions $(Z_M(p,1),Z_M(p,2))$ solve the system of HJB equations (5) and (6). Table 2 shows the explicit solutions when $\sigma_M(p)=1$ with integration constants $D_k(k=1,\dots,3)$. Since the environment is symmetric, the explicit solutions when $\sigma_M(p)=0$ are given by replacing the role of party 1 and 2; for example, $Z_M(p,1)$ for an open interval where $(k_1,k_2,\sigma_M)=(x,y,0)$ is identical to $Z_M(1-p,2)$ for an open interval where $(k_1,k_2,\sigma_M)=(y,x,1)$.

B Omitted Proofs

B.1 Proof of Proposition 1

Suppose $g/s<\alpha_0=\frac{1+2\mu}{1+\mu}$. Then Lemma 1 implies that $p_1^*=1-p_2^*=p^0>1/2$, and that there exists no other threshold belief of party 1 for $p\in[1/2,p^0)$.

Now suppose that there exists a threshold belief of party 1 for some $p\in(0,1/2)$, and let $\tilde{p}_1$ be the smallest such threshold. First, consider the case where $\hat{p}_1\in(1-p^0,1/2)$, so that $k_2(\tilde{p}_1)=0$. Then by the value-matching and the smooth-pasting conditions for $V_1(p)$ at $p=\tilde{p}_1$, from Table 1b we have

$$\begin{aligned} s &= \chi s+\frac{g-\chi s}{\hat{\mu}+1}\tilde{p}_1+C_4\hat{f}(\tilde{p}_1), \\ 0 &= \frac{g-\chi s}{\hat{\mu}+1}-C_4\frac{\tilde{p}_1+\hat{\mu}}{\tilde{p}_1(1-\tilde{p}_1)}\hat{f}(\tilde{p}_1). \end{aligned}$$

Solving for $\tilde{p}_1$ gives $\tilde{p}_1 = \frac{\mu}{g/s-1}$. But for any $g/s < \alpha_0$, it must be that $\tilde{p}_1 > 1$, leading to a contradiction.

Next, consider the case where $\hat{p}_1 < 1 - p^0$, so that $k_2(\tilde{p}_1) = 1$. Again the value-matching and the smooth-pasting conditions for $V_1(p)$ at $p = \tilde{p}$ give

$$\begin{aligned}(1-\chi)s + \chi C_5 f(1-\tilde{p}_1) &= \frac{\mu+1}{\hat{\mu}+1} g\tilde{p}_1 + \frac{\hat{\mu}\chi}{\hat{\mu}+\mu+1} C_5 f(1-\tilde{p}_1) + C_6 \hat{f}(\tilde{p}_1),\\ \chi C_5 f'(1-\tilde{p}_1) &= \frac{\mu+1}{\hat{\mu}+1} g + \frac{\hat{\mu}\chi}{\hat{\mu}+\mu+1} C_5 f'(1-\tilde{p}_1) - C_6 \hat{f}'(\tilde{p}_1).\end{aligned}$$

Combining the above two equations, we have

$$(1+\mu)\chi C_5 f(1-\tilde{p}_1) = -\mu s + (-(1-\chi)s + (1+\mu)g)\tilde{p}_1.$$

On the other hand, the value-matching condition for $W_1(p)$ at $\tilde{p}_1$ gives $s = C_5 f(1-\tilde{p}_1)$. Substituting it to the above equation gives the solution for $\tilde{p}_1$:

$$\tilde{p}_1 = \frac{\mu + (\mu+1)\chi}{(\mu+1)(g/s) - (1-\chi)}.$$

However, whenever $g/s < \alpha_0$ it must be that $\tilde{p}_1 > 1/2$, leading to a contradiction.

B.2 Proof for Propositions 2-3

Suppose $g/s \geq \alpha_0$. Then the single-party threshold p^0 is less than $1/2$, and Lemma 1 implies that $p_1^* \leq 1/2$ and $p_2^* \geq 1/2$. Moreover, from Table 1a we have

$$V_1(p) = \quad gp + C_2 f(p), \qquad \text{for } p > 1/2, \tag{9}$$

$$W_1(p) = \quad C_5 f(1-p), \qquad \text{for } p < 1/2. \tag{10}$$

It remains to characterize the parties' behavior in the unfavorable range of beliefs (that is, $p \leq 1/2$ for party 1 and $p \geq 1/2$ for party 2). In the following subsections, we characterize the equilibria in the following three cases:

1. $p_1^* = p_2^* = 1/2$;

2. $p_1^* < 1/2$ and $p_2^* > 1/2$;

3. $p_1^* = 1/2$ and $p_2^* > 1/2$; or $p_1^* < 1/2$ and $p_2^* = 1/2$.

B.2.1 Case 1: Proposition 3

Consider the case in which $p_1^* = p_2^* = 1/2$. Since both parties play the safe action at $p = 1/2$, we have $V_i(1/2) = W_i(1/2) = s$ for all $i = 1,2$ regardless of the median voter's strategy at $p = 1/2$. These boundary conditions give us the values of integration constraints $C_2 = 2s - g$ and $C_5 = 2s$ in (9) and (10).

Next, observe that party 1's behavior for $p < 1/2$ does not depend on party 2's behavior for $p > 1/2$ (and vice versa). This is because if the prior is less than $1/2$, the posterior never reaches $p \in (1/2, 1)$. Moreover, if $k_1(p)$ is a best response of party 1 for $p < 1/2$, then $k_2(p) = k_1(1-p)$ is a best response of party 2 for $p > 1/2$ by symmetry. Therefore, it remains to show the best response of party 1 for $p < 1/2$ to $k_2(p) = 1$ and $\sigma_M(p) = 0$.

Suppose $p_1^* = 1/2$ is party 1's only threshold point. Then $\sigma_1(p) = S$ for all $p < 1/2$, and by Table 1,

$$V(p) = (1-\chi)s + 2\chi s f(1-p).$$

In order to check the optimality of the profile, recall that $b_1(p;V_1) = \frac{p}{\mu}\{g - V_1(p) - (1-p)V_1'(p)\}$ and $c_1(p) = s - pg$ are the value and the opportunity cost of experimentation for party 1, respectively. Therefore, choosing S for all $p < 1/2$ is party 1's best response if and only if $b(p;V_1) \leq c(p)$ for any $p \in (0,1/2)$, or

$$(1+\mu)2\chi s f(1-p) \geq -\mu s + (-(1-\chi)s + (1+\mu)g)p, \tag{11}$$

for any $p \in (0,1/2)$. A simple calculation shows that (11) is equivalent to

$$\frac{g}{s} < \alpha_1(\mu,\chi) \equiv \inf_{p\in(0,1/2)} \left(\frac{1-\chi}{1+\mu} + \frac{\mu}{(1+\mu)p} + 2\chi H(p) \right), \tag{12}$$

where $H(p) = \left(\frac{p}{1-p}\right)^{\mu}$. To derive the closed-form formula of $\alpha_1(\mu,\chi)$, let p^z be the minimizer of the right-hand side of above equation. Since $H'(p) = \frac{\mu}{p(1-p)}H(p)$, an interior minimizer p^z satisfies

$$\left(\frac{p^z}{1-p^z}\right)^{\mu+1} = \frac{1}{2\chi(1-\mu)}. \tag{13}$$

Since the left-hand side of (13) is strictly increasing from zero to one as p^z moves from zero to 1/2, there exists unique $p_z \in (0,1/2)$ that satisfies (13) if $\chi > \frac{1}{2(1+\mu)}$. If $\chi \leq \frac{1}{2(1+\mu)}$, then the right-hand side of

(12) is strictly decresing in p and thus $p^z = 1/2$. Applying the value of p^z to (12) yields

$$\alpha_1(\mu,\chi) = \begin{cases} \frac{1+2\mu}{1+\mu}(1+\chi) & \text{if } \chi \leq \frac{1}{2(1+\mu)} \\ \frac{1}{p^z} - \frac{\chi}{1+\mu} & \text{if } \chi > \frac{1}{2(1+\mu)} \end{cases}$$

It is straightforward to check that $\alpha_1 > \alpha_0 = \frac{1+2\mu}{1+\mu}$ and that α_1 is increasing in χ when $\chi \leq \frac{1}{2(1+\mu)}$. It remains to check $\alpha_1(\mu,\chi)$ is increasing in χ when $\chi > \frac{1}{2(1+\mu)}$. From (13), the implicit function theorem implies that

$$\frac{\partial p^z}{\partial \chi} = -\frac{p^z(1-p^z)}{2\chi^2(1+\mu)^2}H(p^z).$$

Therefore, we have

$$\begin{aligned} \frac{\partial \alpha_1(\mu,\chi)}{\partial \chi} &= \frac{1}{1+\mu}\left(\frac{1}{2\chi^2(1-\mu)}\left(\frac{1-p^z}{p^z}\right)^{\mu+1} - 1\right) \\ &= \frac{1}{1+\mu}\left(\frac{1}{\chi} - 1\right) > 0, \end{aligned}$$

where the second equation is from (13).

B.2.2 Case 2: Proposition 2

Suppose $p_1^* < 1/2$ and $p_2^* > 1/2$. First we show that for any $p_2 > 1/2$, party 1's best response threshold p_1^* is uniquely determined. The intuition is as follows. Observe that if the prior belief were less than p_2, the posterior never falls into $p \in (p_2, 1)$. Therefore, the party 1's optimal response does not depend on party 2's action for $p \in (p_2, 1)$. Using the fact that party 2 plays the safe action for all $p \leq p_2$, party 1's best response is determined.

Fix any $p_2 > 1/2$, and let $A = \frac{\hat{\mu}}{\hat{\mu}+\mu+1}$ and $B = \frac{\mu+1}{\hat{\mu}+1}$. Then the following five boundary conditions determine the unique p_1^*:

1. Value-matching condition of V_1 at $p = p_1^*$:

$$(1-\chi)s + \chi C_5 f(1-p_1^*) = Bgp^* + A\chi C_5 f(1-p_1^*) + C_6 \hat{f}(p_1^*), \tag{14}$$

2. Smooth-pasting condition of V_1 at $p = p_1^*$:

$$\chi C_5 f'(1-p_1^*) = Bg + A\chi C_5 f'(1-p_1^*) + C_6 \hat{f}'(p_1^*), \tag{15}$$

3. Value-matching condition of V_1 at $p = 1/2$:

$$Bg + A\chi C_5 + C_6 = g + C_2, \tag{16}$$

4. Value-matching condition of W_1 at $p = 1/2$:

$$C_5 = \chi g + A\chi C_2 + C_3, \tag{17}$$

5. Value-matching condition of W_1 at $p = p_2$:

$$\chi g(p_2) + A\chi C_2 f(p_2) + C_3 \hat{f}(1-p_2) = (1-\chi)s + \chi g(p_2) + \chi C_2 f(p_2). \tag{18}$$

Combining (14) and (15) yields

$$(1+\mu)\chi C_5 f(1-p_1^*) = -\mu s + (-(1-\chi)s + (1+\mu)g)p_1^*. \tag{19}$$

Simplifying (14) and (16)-(18) yields

$$C_2 = \frac{\hat{f}(1-p_2)(\hat{f}(p_1^*)(B-1)g + (1-\chi)s - Bgp^*) + \chi\alpha(p_1^*)(\chi\hat{f}(1-p_2)g + (1-\chi)s)}{\hat{f}(p_1^*)\hat{f}(1-p_2) - \chi^2\alpha(p_1^*)\alpha(1-p_2)}, \tag{20}$$

$$C_5 = \frac{\chi\alpha(p_1^*)(\hat{f}(p_1^*)(B-1)g + (1-\chi)s - Bgp^*) + \hat{f}(1-p_2)(\chi\hat{f}(1-p_2)g + (1-\chi)s)}{\hat{f}(p_1^*)\hat{f}(1-p_2) - \chi^2\alpha(p_1^*)\alpha(1-p_2)}, \tag{21}$$

where $\alpha(p) = A\hat{f}(1-p) + (1-A)f(p)$.

Since $\lim_{\chi\to 0} C_5 < \infty$ for any $p_1^* \in (0,1/2)$ and $p_2 \in (1/2,1)$, the left-hand side of (19) converges to zero as $\chi \to 0$. Therefore, for any $p_2 \in (1/2,1)$, the best-response threshold p_1^* (that is, the solution of (19)) is unique for small enough χ, and p_1^* converges to $p^0 = \mu g/(\mu s + (g-s))$ as $\chi \to 0$. Note that $p^0 \in (0,1/2)$ since $g/s > \alpha_0 = \frac{1+2\mu}{1+\mu}$. Therefore, if χ is close to zero, we conclude that there exists a unique pair (p_1^*, p_2^*) $(p_1^* \in (0,1/2), p_2^* \in (1/2,1))$ such that p_i^* is the best response threshold to p_j^*. Furthermore, $p_1^* = 1 - p_2^*$, which implies that the equilibrium is symmetric.

B.3 Proof of Proposition 5

First, let us show that the Markov profiles after deviation of each player constitute Markov perfect equilibria. Since the profile after a deviation by the median voter is identical to one in Proposition 3, we only needs to show that the profile after a deviation by a political party is a MPE. Without loss of generality, let us consider the profile after party 1's deviation, that is, the profile that is similar to one in Proposition 3 but with belief threshold of $1/2+x_2$.

Observe from Table 1 that $V_1(p) = gp + C_2 f(p)$ for $p > 1/2+x_2$ and $W_1(p) = C_5 f(1-p)$ for $p < 1/2+x_2$. Moreover, similar to the profile in Proposition 3, the boundary conditions are given by $V_1(1/2+x_2) = W_1(1/2+x_2) = s$. Therefore, the value of integration constants are

$$C_2 = \frac{s - g\cdot(1/2+x_2)}{f(1/2+x_2)}, \qquad C_5 = \frac{s}{f(1/2-x_2)}.$$

Paty 1's optimality for $p > 1/2+x_2$ is given by Lemma 1. To show party 1's optimality for $p < 1/2+x_2$, note that the value function $V_1(p)$ for such belief range is $V_1(p) = (1-\chi)s + \chi C_5 f(1-p)$. Then playing S is optimal for $p < 1/2+x_2$ if and only if $b(p;V_1) \le c(p)$ for any $p \in (0,1/2)$. A calculation similar to one in Subsection B.2 yields that the optimality condition is equivalent to

$$\frac{g}{s} < \tilde{\alpha}_1(\mu,\chi,x_2) \equiv \inf_{p\in(0,1/2)} \left(\frac{1-\chi}{1+\mu} + \frac{\mu}{(1+\mu)p} + \frac{\chi H(p)}{f(1/2+x_2)} \right). \tag{22}$$

Since $\tilde{\alpha}_1(\mu,\chi,x_2)$ converges to $\alpha_1(\mu,\chi)$ as $x_2 \to 0$, we conclude that for small enough x_2 the above profile is optimal for party 1. A similar anlaysis shows the optimality of party 2.

References

ARAGONES, E., PALFREY, T. and POSTLEWAITE, A. (2007). Political Reputations and Campaign Promises. *Journal of the European Economic Association*, **5** (4), 846–884.

ASHWORTH, S. and DE MESQUITA, E. B. (2008). Electoral Selection, Strategic Challenger Entry, and the Incumbency Advantage. *The Journal of Politics*, **70** (04), 1006.

BANKS, J. S. and SUNDARAM, R. K. (1998). Optimal Retention in Agency Problems. *Journal of Economic Theory*, **82** (2), 293–323.

BESLEY, T. (2004). Joseph Schumpeter Lecture: Paying Politicians: Theory and Evidence. *Journal of the European Economic Association*, **2** (2-3), 193–215.

BOLTON, P. and HARRIS, C. (1999). Strategic Experimentation. *Econometrica*, **67** (2), 349–374.

BONATTI, A. and HÖRNER, J. (2011). Collaborating. *American Economic Review*, **101** (2), 632–663.

CALLANDER, S. and HUMMEL, P. (2014). Preemptive Policy Experimentation. *Econometrica*, **82** (4), 1509–1528.

DIXIT, A. K. and GUL, F. (2000). The Dynamics of Political Compromise. *Journal of Political Economy*, **108** (3), 531–568.

FEREJOHN, J. A. (1986). Incumbent Performance and Electoral Control. *Public Choice*, **50** (1/3), 5–25.

GUL, F. and PESENDORFER, W. (2012). The War of Information. *Review of Economic Studies*, **79** (2), 707–734.

KELLER, R. G., RADY, S. and CRIPPS, M. W. (2005). Strategic Experimentation with Exponential Bandits. *Econometrica*, **73** (1), 39–68.

KLEIN, N. and RADY, S. (2011). Negatively Correlated Bandits. *Review of Economic Studies*, **78** (2), 693–732.

MASKIN, E. and TIROLE, J. (2004). The Politician and the Judge: Accountability in Government. *American Economic Review*, **94** (4), 1034–1054.

SAMUELSON, L. (1987). A Test of the Revealed-Preference Phenomenon in Congressional Elections. *Public Choice*, **54** (2), 141–169.

STRULOVICI, B. (2010). Learning While Voting: Determinants of Collective Experimentation. *Econometrica*, **78** (3), 933–971.