Preemptive Experimentation Under Alternating Political Power*

Steven Callander and Patrick Hummel†

February 6, 2013

Abstract

A fundamental feature of politics is that parties of differing ideologies frequently take turns holding political office. Although normatively appealing, it is well-known that these transitions of power induce a time-inconsistency in policymaking that can significantly distort policy choices when a real state variable (such as a budget deficit) connects periods. In a two period model, we analyze the logic of taking turns when the only linkage across policymaking periods is information. We show that even without an ability to tie-the-hands of his successor in any way, an early office holder can nevertheless shape future policy choices by experimenting with policy and shaping the information that his successor holds. We show that this incentive to *preemptively experiment* emerges for policy issues of moderate complexity, and we explore how it impacts and can *improve* the efficiency of policymaking across time. We relate this behavior to recent episodes in domestic and foreign politics. We also explore an alternative application of the model to the choice of political system itself, and use it to interpret the dynamics of democratization as strategic experimentation and information revelation.

Keywords: Political economy, alternating power, time-inconsistency, policy experimentation, democratization.

---

*For helpful comments we thank Maggie Penn, Dave Baron, Alessandra Casella, Seema Jayachandran, Johannes Hörner, Andrea Prat, Moses Shayo, Ken Shotts, the editor, and the anonymous referees. Davin Raiha provided research assistance.

†Corresponding author: Steven Callander, Graduate School of Business, Stanford University, Stanford, CA 94305; sjc@gsb.stanford.edu.

1 Introduction

A fundamental feature of democracy is that parties of differing ideologies frequently take turns holding political office. So fundamental is this feature that the (peaceful) transition of political power is often taken as evidence of a healthy democracy.

Despite this, the wisdom of parties taking turns in power remains ambiguous. Changing the party of government is costly in both time and resources, and uncertainty over future governments can undermine investment and economic behavior. More perniciously, it is well known that taking turns in political power induces a time-inconsistency in policymaking when a real state variable, such as a budget deficit, connects policymaking periods, and that this time-inconsistency significantly distorts policy choices (Persson and Svensson 1989; Tabellini and Alesina 1990).

In this note we address taking turns in the absence of a real state variable. Instead, we explore the logic of taking turns when the linkage across policymaking periods is purely *informational*.

Information plays an important role in real policymaking. Finding a good policy is a hard problem as policymakers must grapple with nontrivial uncertainty about which policies produce which outcomes. Frequently the only way to learn is to try a policy and live with the outcome. This difficulty does, however, provide the early office holder with an opportunity. By experimenting with policy, the early office holder can shape the choice of his successor by shaping the information that he holds. We show that on issues of moderate complexity this incentive drives equilibrium behavior. Thus, even without the ability to tie-the-hands of his successor in any formal way, an intertemporal linkage emerges in equilibrium that is purely informational. We explore how this incentive impacts and can improve the quality of policymaking across time.

The key idea is the notion of a *preemptive* policy experiment. Consider a situation in which the status quo policy favors the early incumbent and policymaking is over two periods. If the early incumbent's control on power were permanent, his desire to experiment would be small to nonexistent. However, if his control on power is temporary, and he is to be succeeded by a party with different ideology, he may face the prospect of a large, future experiment that could move the outcome well away from his ideal. We show how the early incumbent, anticipating his successor's reaction, preemptively engages in a small policy experiment in the hope of heading off a large subsequent experiment. As a result causality runs both ways through time: The preemptive experiment shapes the choice of a future office holder, and at the same time the looming shadow of a successor with different preferences shapes the choices of the early office holder.

In many situations the preemptive policy experiments that are undertaken are socially efficient. By preemptively experimenting toward his political opponent, the early incumbent moderates policy and potentially avoids a large experiment. To a representative median voter situated between the parties, this preemptive experimentation is the moderation that she seeks. Strikingly, this moderation occurs even with a fixed, predetermined sequence of political control and without the disciplining effect of elections. In fact, we show that in this situation, the median voter is strictly *better off* alternating power between the parties with policy uncertainty than she would be if the mapping from policies to outcomes were fully revealed and all policy uncertainty removed. It is still the case here that policy choice is time-inconsistent, yet this result shows that when the linkage across periods is informational rather than via a real state variable, the time-inconsistency can work to the benefit rather than the detriment of voters.

Although preemptive policy experiments can be socially beneficial, the individual motivation behind them is not. In fact, the logic of preemptive experimentation is pure Machiavellianism: The early incumbent often hopes that his preemptive experiment *fails* rather than succeeds. A policy failure is initially costly, yet it is valuable as it dissuades the successor government from even bolder policy experimentation. In many situations, such as described above, this Machiavellian logic works to the social benefit. However, this isn't always the case and we identify the possibility for preemptive experiments to undermine social welfare and induce Pareto inefficient policy choices.

The premise of the model is that policy making is often difficult yet not entirely random. Mistakes are frequently made and intended outcomes are difficult to achieve. In fact, outcomes are often the opposite of that intended, as captured in Merton's (1936) famous Law of Unintended Consequences. Nevertheless, policy choice is not random and it is possible to learn from experience, from both successes and failures. To capture this richness I employ the formulation of uncertainty introduced in Callander (2011, 2011b).[1] In this formulation, the policy space is the real line and policies are mapped into outcomes by a function, where this function is represented as the realized path of a Brownian motion. Policymakers do not know the realized path, but are aided by the structural knowledge of how the path is generated (the parameters of the Brownian motion). This knowledge, combined with the practical lessons of experience as it accumulates, guides behavior and enables policymakers to do more than simply randomize with their policy choices. Importantly, this formulation provides a notion of the size and direction of policy experimentation that allows me to capture the distinction between small and large experiments that is essential to the notion of a preemptive experiment.

To expose most clearly the logic of taking turns, the model of strategic behavior we build is as parsimonious as possible. We limit attention to two periods with two players and focus on the intertemporal incentives to actively experiment when control of power is temporary and the order of control predetermined. Moreover, throughout the analysis we intermingle examples with general results to convey most clearly the intuition and robustness of the underlying phenomena. In the discussion section we describe several real-world examples of preemptive policy experiments from foreign policy and a recent episode from domestic politics in Israel.

[1] Callander (2011) models uncertainty in product markets and Callander (2011b) adapts the specification to policy making under uncertainty. Both of those papers exclusively use the Brownian motion, whereas here we employ both the Brownian motion and the discrete random walk, amongst other differences.

Throughout the paper we frame the model in the language of politics, specifically the liberal-conservative dimension of modern politics. However, broader interpretations are available.[2] One intriguing possibility is for the choice to be over the political system itself. Two patterns that emerge in the history of democratization are (i) that democracy often fails and countries revert to dictatorship, and (ii) that the evolution of democracy is often iterative, with democratic freedoms progressively added (such as in the U.K.). The model offers a novel twist on these phenomena by suggesting they emerge from experimentation and learning about the efficacy of democracy itself. We return to this interpretation in the discussion section and consider how preemptive experimentation is consistent with strategic democratization and the backsliding from democracy to dictatorship.

1.1 Related Literature

The question of taking turns in political office has recently been addressed forcefully by Acemoglu, Golosov, and Tsyvinski (2011). In their model, power fluctuates stochastically between an arbitrary number of parties over an infinite horizon. They provide a positive result: That more frequent power fluctuations increases cooperation and more effectively smooths consumption across time. My paper shares with Acemoglu *et al* an emphasis on how the outcomes an agent receives when outside office impact their choices within office. The models and intuitions are otherwise distinct. The most important distinction is the planning horizon: In Acemoglu *et al* the infinite horizon is essential to maintaining the cooperative equilibrium, whereas we analyze a two period model that cannot sustain such cooperative behavior in equilibrium.

The restriction to two periods is an obvious simplification relative to the recent literature on dynamic political economy, of which Acemoglu *et al* is one example. Nevertheless, in no paper in that literature is information the intertemporal connection and strategic experimentation plays no role. Moreover, capturing policy experimentation in a realistic way – with an ideological space and a conception of distance and direction of policy choice – requires a rich model of uncertainty and learning that complicates the stage-game. As will become evident as the analysis proceeds, the idea of a small preemptive experiment heading off a larger experiment is appealingly simple and intuitive, yet the supporting logic is subtle. For it to emerge in equilibrium it is necessary that policy outcomes are correlated and that policy choice not be binary. These properties are exactly captured by the Brownian motion representation of uncertainty, and are not present in classic, simpler models of policymaking under uncertainty.

Formally, the Brownian motion representation corresponds to a bandit model with an infinite number of correlated, deterministic bandits. Working with correlated bandits is

[2]At its heart, democratic politics is a problem of control rights and the allocation of control rights is important in any number of applications outside of political economy. Examples of taking turns include a manager ceding control mid-project to his boss, a joint venture partner controlling an investment early in a project's life before handing it to the other partner, or two parents taking turns deciding how to discipline their child.

notoriously difficult. The trade-off made here is to incorporate this stage-game richness at the expense of only two periods of policy making, an approach that resonates with that of Persson and Svensson (1989) and Tabellini and Alesina (1990) in the first generation of dynamic political economy models. That said, the Brownian motion representation offers enough mathematical structure that future work may very well be able to incorporate the two-period insights into fully dynamic models. The planning horizon distinguishes this paper from my previous work (Callander 2011, 2011b). In those papers, decision making is entirely myopic and experimentation passive. Although the horizon here is only two periods, active experimentation and the expectation of a future payoff from a utility-lowering experiment today is at the heart of the results. This paper is also distinct in that control is shared and preferences are not common across decision makers.

The paper also speaks to a large literature in political science on whether and how a legislature is able to tie-the-hands of its successors. The U.S. constitution explicitly puts every Congress on equal footing, with each able to undo what another does (Article I, Section 7). This has led to a large literature on how institutional design can be used to *insulate* policy choices, with a focus on bureaucratic design features that insulate policies at the cost of efficiency (bureaucracies are "inefficient by design"; Moe 1989). A challenge to this argument is that the insulating mechanisms can themselves be overturned and evidence to this effect is provided by Berry, Burden, and Howell (2010). My model speaks to this literature by side-stepping the possibility of formal constraints on successor governments. Instead we show how an early Congress can impact its successors simply by conveying information to them through policy experiments, such as is evident in how the New Deal (a success) and prohibition (a failure) shaped subsequent policy debates throughout the 20^{th} century.

This line of reasoning sidesteps the commitment problem by rendering it irrelevant. Once information is revealed it cannot be forgotten (although it can be debated; such as is the case with the New Deal).[3] The irreversibility of information revelation is also of interest as it generates a policy dynamic that is path dependent, and a path dependence that does not require formal institutional constraints. Path dependence is a concept that has proved popular in the qualitative political science literature (Pierson 2000, 2004) but has yet to yield substantially to formal reasoning (Page 2006).[4]

2 The Model

In each of two periods, $t = 1, 2$, a policy is chosen by the incumbent office holder, who is either the Democrat (D) or the Republican (R). The chosen policy produces an outcome according to the function, ψ, where both the policy and outcome spaces are single-dimensional.

[3]The inability to commit is also at the heart of the democratization literature.

[4]Path dependence is also highlighted by Gibbons (2010) as an important area in organizational economics that is rich in evidence but short on theory. Surprisingly, the organizational economics literature on control rights has not focused on taking turns. The closest papers study shared control in one-shot models with an emphasis on communication (Alonso, Dessein, Matouschek 2008; Alonso 2008).

The politicians and voters care about outcomes and only indirectly about policies. The ideal outcomes for the Democrat and for the Republican are $-\gamma$ and $+\gamma$, respectively, with $\gamma > 0$. A representative voter, m, has an ideal outcome of 0. This ordering captures the classic situation of a median voter facing ideological parties arrayed to her either side. For simplicity, we assume each player has quadratic-loss preferences over outcomes.[5] To avoid conflating time-preference with strategic behavior, utility is not discounted across time. Thus, for the pair of policies $\{p_1, p_2\}$, total utility for the Republican is:

$$U_R(p_1, p_2) = -(\psi(p_1) - \gamma)^2 - (\psi(p_2) - \gamma)^2 .$$

The challenge of policymaking is that the mapping from policies to outcomes is not perfectly known. The true mapping, ψ, is chosen by Nature prior to the first period and we model ψ as the realized path of a random walk governed by Brownian motion. The players know the parameters of the motion and seek to learn the realized path of the random walk.[6] The players begin play with knowledge of one point in the mapping, the status quo. This can be thought of as the policy in place at the beginning of play; I normalize the status quo to policy 0 such that its outcome is $\psi(0)$.

To focus on learning, the same policy mapping is in effect for both periods. Changing policy is costless and every policy is equally accessible in both periods. Therefore, the first period incumbent cannot tie the hands of his successor. The outcomes of policies are observed (and experienced) perfectly by all citizens. If a previously untried policy is chosen at any point we say it is a policy *experiment*.[7] If the first period policy is experimental, players learn a second point in the mapping and update their beliefs accordingly.

The policy space and the outcome space are both the entire real line, such that $\psi : \mathbb{R} \to \mathbb{R}$. The Brownian motion is a two parameter process, with the drift, μ, measuring the expected rate of change and the variance, σ^2, the noisiness of the process. Without loss of generality, set $\mu \geq 0$. The left-side panel of Figure 1 depicts one possible realized path of the policy mapping that passes through $(0, \psi(0))$.

For all untried policies, beliefs are uncertain and distributed normally. Suppose the right-most (left-most) known point is $\hat{p}$. For all policies to the right (left) of $\hat{p}$, beliefs depend only on the outcome $\psi(\hat{p})$ and the parameters of the process (by the Markov property). Specifically, the expected outcome for all $p > \hat{p}$ ($p < \hat{p}$) is given by the straight line of slope μ that passes through $(\hat{p}, \psi(\hat{p}))$ and the variance is a linear function of the difference between

[5]The main intuition – and many of the results – generalize under reasonable conditions on utility functions, although the details clutter the presentation. The high degree of tractability of quadratic utility when combined with the normal distribution offers a degree of clarity that rewards the specialization. Risk aversion plays a prominent role in the analysis although, as is well known, risk aversion is present with linear and even convex utility curves when players have an with an internal optimum.

[6]Therefore, they are not trying to learn the drift and/or variance.

[7]This is the standard use of the term: an experiment is the choice of an option with unknown outcome. Note, however, that with deterministic bandits, an experiment in the first period becomes a safe choice in the second.

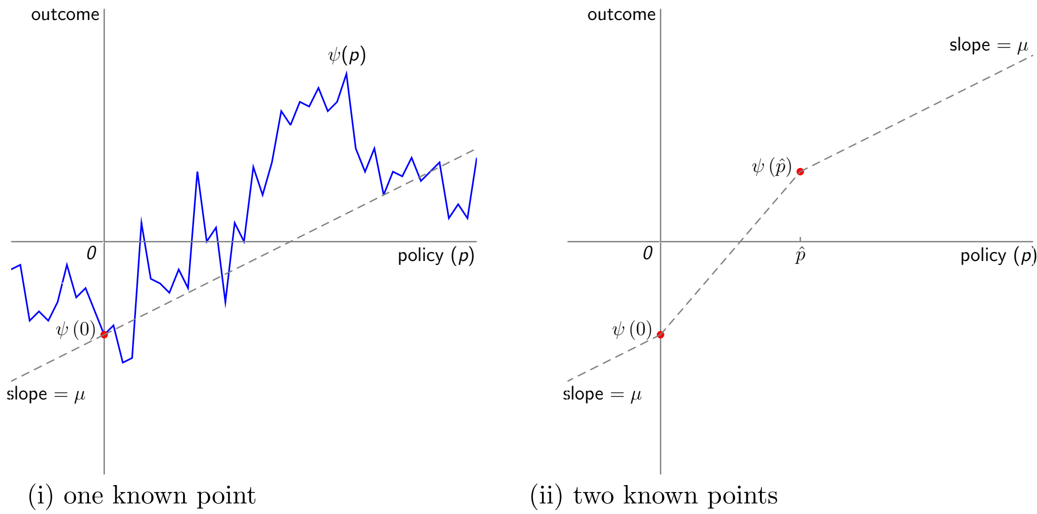

Figure 1: Continuous Policy Mapping: the Brownian Motion

p and $\hat{p}$. Stated formally:

$$\text{Expected Outcome:} \qquad E\psi(p) = \psi(\hat{p}) + \mu(p - \hat{p}), \tag{1}$$

$$\text{Variance:} \qquad var(\psi(p)) = |p - \hat{p}| \, \sigma^2. \tag{2}$$

As beliefs in this range are anchored by only one point I say that they are *open-ended.*

In the first period only a single point in the mapping is known and beliefs are open-ended to either side of 0. The dashed line in the left-side panel of Figure 1 depicts the expected outcome for all policies $p \in \mathbb{R}$.

If the first period incumbent experiments with policy then a second point in the mapping is revealed. On the flanks of the two known points beliefs are again open-ended. This is depicted in the right-side panel of Figure 1, with the drift of line of slope μ extending to the right of $(\hat{p}, \psi(\hat{p}))$ and the left of $(0, \psi(0))$.

Between the two known points beliefs are not open-ended. In this case a Brownian bridge forms and beliefs depend on the nearest known point in either direction. On the bridge beliefs are generated by an intuitive interpolation between the ends. The expected outcomes are given by the straight line between the ends of the bridge. The variance is concave across the bridge, reaching a peak halfway between the ends and equalling zero at the ends. In the analysis to follow one of the known points must be the status quo; thus, we state the beliefs for a bridge between policy 0 and an arbitrary policy $\hat{p}$:

$$\text{Expected Outcome} \quad : \quad E[\psi(p)] = \psi(0) + \frac{p}{\hat{p}}(\psi(\hat{p}) - \psi(0)),$$

$$\text{Variance} \quad : \quad var(\psi(p)) = \left| \frac{p(\hat{p} - p)}{\hat{p}} \right| \sigma^2.$$

A bridge is depicted in the right panel of Figure 1 for $\hat{p} > 0$.

The random walk representation of the policy mapping captures several key features of policymaking and decision making generally:[8]

∘ Expected vs. actual outcomes. Citizens can order policies according to *expected* outcomes but not according to *realized* outcomes. That is, they know which policies are *more likely* to deliver conservative or liberal outcomes but they do not know which policies *do* deliver these outcomes. This represents a natural generalization to uncertain environments of the classic left-right conception of policy.

∘ Partial invertibility. Citizens can learn across policies as the policy path is *partially invertible.* Observing outcomes reveals some information about untried policies but not everything, just as in practice a policy experiment informs future choice without rendering it trivial.

∘ Proportional invertibility. The accuracy of beliefs is decreasing in the distance an untried policy is from a policy for which the outcome is known. This *proportional invertibility* captures the intuition of Lindblom (1959) that greater uncertainty is incurred the more policy is moved from what is known. Proportional invertibility also implies that an experiment reveals more information (in terms of lowering variance) the more novel it is.

∘ Law of Unintended Consequences. In trying to make outcomes better, changes to policy may actually make things worse. Outcomes can overshoot their target or move in the wrong direction, capturing the original insight of Merton (1936).

∘ Issue complexity. The parameters of the stochastic processes provide a simple measure of the *complexity* of the underlying policy issue. The parameters determine how predictable an issue is and how much citizens learn from observing outcomes. We measure complexity by the ratio $\frac{\sigma^2}{|\mu|}$. For low complex issues the outcome of a policy change can be predicted accurately, whereas for more complex issues the outcome of a policy change is just as likely to become more liberal as more conservative. This allows us to obtain insight into how experimentation and learning vary across simple and complex issues.

The focus of this paper is on the intertemporal incentives of policy makers, and as such we abstract away from voter behavior. Formally, we follow Persson and Svensson (1989) and directly impose particular sequences of political power.[9] To focus on how taking turns affects policy choices, we analyze behavior when the Democrat is in power in the first period and succeeded by the Republican and compare it to when the Democrat holds power continuously; these sequences are denoted $D-R$ and $D-D$, respectively.

For pedagogical reasons we make a simplifying assumption on the status quo point. With only a single known point in the policy mapping, the Democrat is likely to experiment in the first period whether his hold on power is permanent or temporary. To isolate the impact of the incentive to experiment preemptively, we focus on the case where $\psi(0) = -\gamma$, and the known status quo point delivers D's ideal outcome. The usefulness of this restriction is that D will not experiment on his own account if the sequence is $D-D$ and his hold

[8] See Callander (2011) for a more complete account of the properties of the random walk representation.

[9] This is a common modeling device. Other political economy models that similarly assume the transition probabililty to be independent of the incumbent's policy choice include Tabellini and Alesina (1990), Dixit, Grossman, and Gul (2000), and Acemoglu, Golosov, and Tsyvinski (2011).

on power is permanent. Therefore, any experimentation undertaken by D when his hold on power is temporary must be preemptive and aimed at strategically teaching his successor. Although this assumption is extreme, the feature that is substantively important is that the Democrat is less inclined to experiment than is the Republican. Such a relative willingness to experiment holds whenever $\psi(0) < -\gamma$; the comparison is reversed for $\psi(0) > \gamma$, although then the results are simply mirrored in the sequence $R - D$. We also analyze the alternative case of $\psi(0) \in (-\gamma, \gamma)$ and demonstrate that taking turns again affects behavior, although this time it has the reverse effect of dampening experimentation relative to when power is held continuously.

In addition to its realism, the random walk specification is appealing due to its tractability. And this tractability is particularly powerful when combined with quadratic utility. In each period, R's expected utility for policy p reduces to the simple mean-variance form, $-(E\psi(p) - \gamma)^2 - var(\psi(p))$, that combines concisely with the moment expressions above. In addition to its tractability, this form is pedagogically helpful in that it isolates strategic considerations from risk tolerance. As variance enters utility linearly, the trade-off between risk and return in experimentation is independent of the location of the status quo. This independence proves particularly useful with the smoothness of the Brownian motion.

Finally, I introduce some terminology. We say that an experiment *fails* if the outcome moves in the opposite direction to that intended or doesn't change, and that the experiment is a *success* otherwise. We define a *policy program* as a pair of policies for both periods of the model; thus, for every information set, a player possesses an *ideal policy program* that maximizes his/her utility. We denote equilibrium policy choices by an asterisk.

3 Results

We begin with the simple benchmark when uncertainty is removed and knowledge of the policy mapping complete. Behavior here takes a simple form. With probability one the realized mapping spans the outcome space and when it does each player can identify the policy that delivers his ideal outcome. As the second period incumbent's behavior is then predetermined, the first period incumbent does not attempt to influence his choice. Instead the first period incumbent maximizes his immediate utility and chooses the policy that delivers his own ideal outcome. Information and teaching thus play no role in policy choice. Moreover, the action of an incumbent in any period is independent of the sequence of political power, and for no sequence does the voter receive her ideal outcome of 0 in either period. This result holds for any value of $\psi(0)$.

Lemma 1 *For arbitrary $\psi(0)$ and known policy mapping ψ, the voter does not get her ideal policy program in any sequence of political power with probability one.*

The simplicity of Lemma 1 is illuminating. It establishes that the sequence of control has *no* impact on policy choice when there is complete knowledge of how policy maps into outcomes. This sets a benchmark that contrasts sharply with behavior when there is uncertainty about the consequences of policy.

Behavior when there is uncertainty about the consequences of policy builds off a simple lemma describing the Republican's second period policy choice should the Democrat not experiment in the first period.[10]

Lemma 2 *Suppose $p_1 = 0$ for sequence $D - R$. The second period policy choice is:*
(i) $p_2^* = 0$ if $\gamma \leq \frac{\sigma^2}{4\mu}$.
(ii) $p_2^* > 0$ such that $E\psi(p_2^*) = -\gamma + \mu p_2^* = \gamma - \frac{\sigma^2}{2\mu}$ if $\gamma > \frac{\sigma^2}{4\mu}$.

With this lemma in hand, first period behavior for low polarization is immediate. If $\gamma \leq \frac{\sigma^2}{4\mu}$ then R will not experiment in the second period if D does not. With nothing to preempt, the Democrat need not experiment in the first period. For low polarization, therefore, taking turns has no impact on first period policy choice.

For moderate levels of polarization, however, D may have an incentive to preemptively experiment to influence R's second period policy choices. Complete equilibrium behavior for D in the first period is characterized by the following proposition:

Proposition 1 *In equilibrium, D follows a strategy characterized by two cutoffs, γ^* and γ^{**} satisfying $\frac{\sigma^2}{4\mu} < \gamma^* < \gamma^{**}$ such that the following properties are satisfied:*

(1) D chooses the policy $p_1 = 0$ if $\gamma \leq \frac{\sigma^2}{4\mu}$ or $\gamma > \gamma^{**}$.
(2) D chooses a policy $p_1 > 0$ that varies continuously with γ for values of $\gamma \in (\frac{\sigma^2}{4\mu}, \gamma^*)$.
(3) D chooses a policy $p_1 < 0$ that varies continuously with γ for values of $\gamma \in (\gamma^*, \gamma^{**})$.
(4) D either chooses a policy $p_1 > 0$ or a policy $p_1 < 0$ when $\gamma = \gamma^*$.
(5) D either chooses a policy $p_1 < 0$ or the policy $p_1 = 0$ when $\gamma = \gamma^{**}$.

Before describing the intuition for why equilibrium takes the form given in Proposition 1, it is first helpful to understand why experimentation can be beneficial for D. The short-term impact of this experiment is unambiguously negative for the Democrat: He can ensure his ideal outcome by not experimenting, yet by experimenting he guarantees the outcome will not be at his ideal. The preemptive experiment is profitable to the Democrat as it changes the Republican's behavior in the second period in a favorable way and this benefit can outweigh the first period cost.

In particular, if D experiments with a policy $p_1 > 0$, there is a chance that this experiment will fail and deliver an outcome that is even further from R's ideal point than the original outcome. This would induce R to use an even larger experiment in the second period if R were obligated to continue to experiment. But R is not obligated to experiment further and, under some conditions, R's optimal response is to backtrack to the policy 0 and obtain the certain outcome of $-\gamma$. Although this outcome is not R's ideal, it delivers a risk-free outcome that is preferable to further experimentation from a distant starting point. In this

[10] This is an adapted version of Proposition 1 in Callander (2011).

event we say that the Republican gets "stuck" as he would like to move the outcome to the right – and he knows in which direction he should move policy to make that most likely – yet he learns that finding a good policy is harder than he had thought and abandons the search, getting stuck at a less desirable outcome.

This desire to stop R from experimenting in the second period can also induce D to experiment with a policy $p_1 < 0$ in the first period. Although the movement of policy to the left is Pareto inefficient, the logic of the preemptive experiment is the same as when the experiment is to the right. In both cases the Democrat is attempting to get his successor stuck so that he doesn't experiment with a distant policy. To understand why D may find it optimal to experiment to the left, note that a surprising result of an experiment to the left (an outcome that moves up) can create a local maximum in R's expected utility that causes R to give up on experimenting further. This may in turn result in a more favorable second period policy outcome for D than if D had not experimented in the first period. The only difference between the preemptive experiments to the right and the left is that an experiment p_1 to the left hopes to get R stuck at $p_2 = p_1$ whereas an experiment to the right hopes to get R stuck at policy $p_2 = 0$.

With this background in mind, we now present the intuition for Proposition 1. First consider what happens when γ is only slightly greater than $\frac{\sigma^2}{4\mu}$. When γ is only slightly greater than $\frac{\sigma^2}{4\mu}$, even making a tiny deviation from the policy $p_1 = 0$ will result in R getting stuck in the second period with probability close to $\frac{1}{2}$. If γ is only slightly greater than $\frac{\sigma^2}{4\mu}$ and D chooses a policy p_1 that is only slightly greater than zero in the first period, then R will get stuck in the second period after almost any second period policy outcome $\psi(p_1) < -\gamma$. And if D chooses a policy p_1 that is only slightly less than zero in the first period, then R will get stuck in the second period after almost any second period policy outcome $\psi(p_1) > -\gamma$. These events both happen with probability close to $\frac{1}{2}$ since the expected policy outcome $\psi(p_1)$ only differs from $-\gamma$ by an amount on the order of $|p_1|$, whereas the standard deviation in this policy outcome is on the order of $\sqrt{|p_1|}$, which is much greater than $|p_1|$ for small values of $|p_1|$.

Thus when γ is only slightly greater than $\frac{\sigma^2}{4\mu}$, D can make it very likely that R will get stuck in the second period at a very small cost in terms of expected first period policy outcomes. Since the benefits from making R get stuck will turn out to be an order of magnitude greater than these small costs, D will choose a policy $p_1 \neq 0$ when γ is only slightly greater than $\frac{\sigma^2}{4\mu}$.

Next consider what happens when γ is significantly larger than $\frac{\sigma^2}{4\mu}$. When γ is significantly larger than $\frac{\sigma^2}{4\mu}$, it is extremely unlikely that R will get stuck. R's ideal point in these cases is so far from the status quo or the types of policies that are likely to arise if D makes a moderate deviation from the status quo in the first period that it is almost certain that R will choose to experiment in the second period even if D chooses a policy $p_1 \neq 0$. Thus choosing a policy $p_1 \neq 0$ in the first period ensures that D will obtain a first period policy that is less favorable in expectation than the policy that would be obtained by choosing

$p_1 = 0$ while giving D little chance of obtaining a more favorable second period policy as a result of R getting stuck. As a result, D will always choose the policy $p_1 = 0$ for large values of γ.

Furthermore, it is generally the case that when γ becomes larger relative to $\frac{\sigma^2}{4\mu}$, R's ideal point becomes further from both the status quo and the types of policies that are likely to arise if D makes a moderate deviation from the status quo in the first period. As a result, it becomes harder to stop R from experimenting by choosing a policy $p_1 \neq 0$ when γ becomes larger. Furthermore, the size of the deviation p_1 that D must make in order to obtain a significant probability of stopping R from experimenting in the second period becomes larger as γ becomes larger relative to $\frac{\sigma^2}{4\mu}$. Thus when γ becomes larger relative to $\frac{\sigma^2}{4\mu}$, the costs from trying to make R get stuck in the second period become larger, and the probability of successfully making R get stuck becomes smaller. Thus in general, preemptive experimentation is less attractive for D in the first period for larger values of γ, and there will be some cutoff γ^* such that D preemptively experiments in the first period when $\gamma \in (\frac{\sigma^2}{4\mu}, \gamma^*)$ but not when $\gamma > \gamma^*$.

Now consider the question of whether D would prefer to choose a policy $p_1 > 0$ or $p_1 < 0$ for values of $\gamma > \frac{\sigma^2}{4\mu}$ where either the optimal policy $p_1 > 0$ or the optimal policy $p_1 < 0$ could be preferred to the policy $p_1 = 0$. Note that D's decision as to whether he should choose a policy $p_1 > 0$ or a policy $p_1 < 0$ is driven by two competing factors. On one hand, a policy $p_1 > 0$ is more appealing than a policy $p_1 < 0$ because D will obtain a better second period policy in the circumstances where R gets stuck if he chooses a policy $p_1 > 0$ than if he chooses a policy $p_1 < 0$. If R gets stuck after D chooses a policy $p_1 > 0$, then R will choose the policy $p_2 = 0$ and D will necessarily obtain his preferred second period policy outcome. But if R gets stuck after D chooses a policy $p_1 < 0$, then R will choose the policy $p_2 = p_1$, and the second period policy outcome will satisfy $\psi(p_2) > -\gamma$, meaning the second period policy outcome will not quite match D's ideal point.

On the other hand, it is more likely that R will get stuck if D chooses a policy $p_1 < 0$ than it is if D chooses a policy $p_1 > 0$. When D chooses a policy $p_1 > 0$, R will only get stuck if the first period policy outcome is so much lower than $\psi(0) = -\gamma$ that R would prefer the second period policy outcome $\psi(0) = -\gamma$ over the uncertain outcome that would result from choosing a policy $p_2 > p_1$. Since the policy outcome $\psi(0) = -\gamma$ is a significant distance from R's ideal point, this is a relatively difficult condition to meet. However, when D chooses a policy $p_1 < 0$, R gets stuck in the circumstances where the policy outcome $\psi(p_1)$ ends up being sufficiently close to γ that R ends up preferring this policy $\psi(p_1)$ that is significantly greater than $\psi(0) = -\gamma$ over the uncertain outcome that would result from choosing a policy $p_2 > 0$. Since it is easier to get R to accept a policy that is significantly greater than $\psi(0) = -\gamma$ (and thus relatively close to R's ideal point of γ) than it is to get R to accept the policy $\psi(0) = -\gamma$ over some uncertain outcome, this implies that there is a greater chance that R will get stuck if D chooses a policy $p_1 < 0$ than if D chooses a policy $p_1 > 0$.

How D resolves this tradeoff depends on how large γ is compared to $\frac{\sigma^2}{4\mu}$. When γ is only slightly greater than $\frac{\sigma^2}{4\mu}$, then the probability that R gets stuck is very close to $\frac{1}{2}$ regardless of whether D chooses a policy $p_1 > 0$ or $p_1 < 0$. Since the decision as to whether to choose a policy $p_1 > 0$ or $p_1 < 0$ has little effect on the probability that R gets stuck, it is in D's interest to choose the policy that will afford D the greatest utility under the circumstances that R gets stuck, and D will choose a policy $p_1 > 0$ when γ is only slightly greater than $\frac{\sigma^2}{4\mu}$.

But as γ becomes significantly larger than $\frac{\sigma^2}{4\mu}$, the probability that R will get stuck if D chooses a policy $p_1 > 0$ becomes much smaller than the probability that R will get stuck if D chooses a policy $p_1 < 0$. As a result, when γ becomes large enough, it becomes more important for D to choose an action that will result in a greater probability of R getting stuck, and D will prefer a policy $p_1 < 0$ over policies $p_1 > 0$ when γ is significantly larger than $\frac{\sigma^2}{4\mu}$. This explains why equilibrium behavior is in the form given in the proposition.

The intuition for this result emphasizes the importance of R getting stuck in driving D's incentives to experiment in the first period. To further emphasize the importance of getting stuck, we analyze a restricted version of the game in which, by fiat, we exclude the possibility of getting stuck. We refer to the restricted game as the No Stuck Game. Although intuitively simple, some care must be taken in defining what it means to get stuck.[11] The following definition accommodates the possibilities.

The No Stuck Game: Restrict the strategy space as follows: (i) if $p_1 > 0$ and $\psi(p_1) \leq -\gamma$, then $p_2 \in [p_1, \infty)$, and (ii) if $p_1 < 0$ then $p_2 \in [0, \infty)$.

In the No Stuck Game, D cannot benefit from policymaking getting stuck. Therefore, the only advantage to D of a preemptive experiment is if it preempts a larger second period experiment by R. In this setting, the smoothness of quadratic utility and continuous policy deliver an unambiguous verdict: Preemptive experimentation is never optimal.

Proposition 2 *For sequence $D - R$ in the No Stuck Game, $p_1^* = 0$ for all degrees of polarization and issue complexity.*

This result establishes that the potential to get a successor stuck is necessary for the logic of preemptive experimentation to work. Without getting stuck, a preemptive experiment merely imposes a cost on the first period incumbent. The key to proving the result is part (ii) of Lemma 2, the essence of which is that R's second period experiment will take the expected outcome to the same point, regardless of the realization of D's first period experiment.

[11] For preemptive experiments to the right, getting stuck involves R backsliding toward 0 with his policy choice. However, getting stuck cannot be defined simply as R moving back to the left. It is possible that the preemptive experiment is to the left and getting stuck involves R not changing policy, and it is also possible in the continuous model for the outcome of D's experiment to overshoot R's ideal outcome. In this event R backtracks onto the bridge that is formed to get close to his ideal outcome in expectation, although in no sense does the backtracking imply that R is stuck. Another contingency to allow for is experiments arbitrarily close to an excluded point; thus, we cannot merely exclude p_1^*, say, following $p_1^* < 0$ as the neighborhood of p_1^* consistutes equivalent behavior delivering arbitrarily close utility.

Our analysis so far has focused on cases in which the status quo policy outcome is more favorable to D than to R. Although this is a natural state of affairs, it is possible that the status quo produces a more centrist outcome that is relatively closer to R's ideal point. We now analyze the alternative case where the status quo outcome lies closer to R's ideal point than in the above analysis. Two interesting conclusions emerge from this case: First, alternating power again impacts behavior relative to when power is continuously held, and, second, the effect on behavior is reversed relative to the asymmetric case.

To demonstrate this effect, we set the status quo outcome to be $\psi(0) = \gamma - \frac{\sigma^2}{4\mu}$. At this value, the Republican will not experiment with a one-period horizon whereas the Democrat will experiment as long as $\gamma > \frac{\sigma^2}{4\mu}$. However, just as in the case where the status quo satisfied $\psi(0) = -\gamma$, if D experiments in the first period, then R may change his behavior in the second period. Proposition 3 explores how this impacts D's willingness to experiment as polarization varies.

Proposition 3 *Suppose $\psi(0) = \gamma - \frac{\sigma^2}{4\mu}$. A $\gamma' > \frac{\sigma^2}{4\mu}$ exists such that for $\gamma \in [0, \gamma')$, the Democrat doesn't experiment in the first period,* i.e., $p_1^* = 0$. *As $\gamma \to \infty$, the Democrat's choice approaches his one-period optimal policy.*

For moderate levels polarization taking turns affects the Democrat's incentives to experiment in the first period by making the Democrat decide against experimenting for a strictly larger set of degrees of polarization than if power were held continuously. This effect dissipates, however, as polarization increases, as the Democrat's behavior approaches his behavior in the one-period optimal policy when polarization grows.

4 Examples

In this section we illustrate the logic of some of the forces governing strategic experimentation through some simple examples with a discrete policy space. The policy space and the outcome space are both the set of integers, such that $\psi : \mathbb{Z} \to \mathbb{Z}$. The probability of "stepping up" with each unit change in policy – the bias of the process – is q, and without loss of generality set $q \in \left[\frac{1}{2}, 1\right]$; the probability of "stepping down" is then $1 - q$. The left panel of Figure 2 depicts a possible realization of the path that passes through $(p, \psi(p)) = (0, -1)$.

Beliefs over untried policies have the same structure as the continuous case, although they are distributed binomially rather than normally. For policy $p > 0$ in the first period, we can think of there having been p Bernoulli random draws where the mean number of successes is pq. As a "successful" trial increases the outcome by one and a "failure" decreases it by one, the change in expected outcome from the status quo is: $-p + 2pq$.[12] This gives an expected

[12]Start by thinking of all trials as failures to generate the $-p$ term in Equation 3. Each success increases the outcome by 2 and with pq successes in expectation the expression follows.

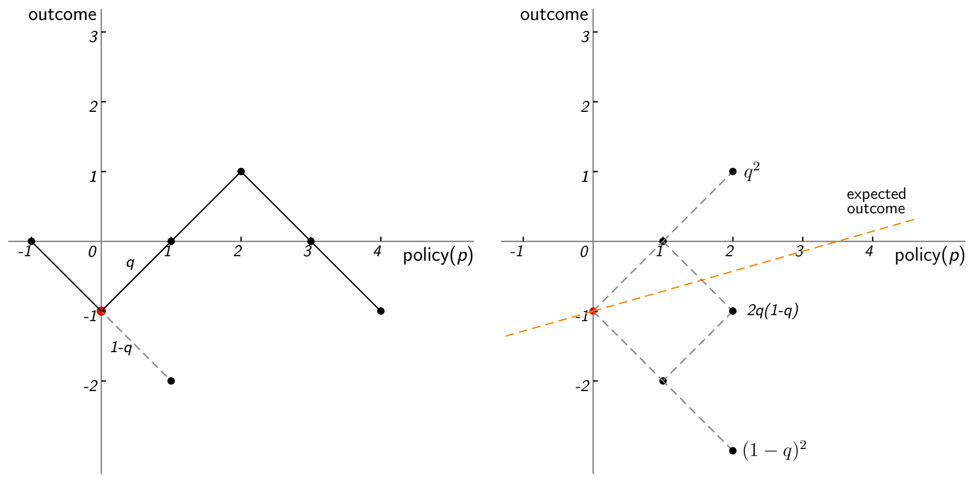

Figure 2: Random walk of step size one (with known point $(0,-1)$).

outcome for $p > 0$ of:

$$\begin{aligned} E\psi(p) &= \psi(0) - p + 2pq \\ &= \psi(0) + p(2q-1). \end{aligned} \tag{3}$$

For policies to the left of the status quo, $p < 0$, the bias q now gives the probability of stepping down. It is easy to see that Equation 3 holds, therefore, for all $p < 0$ as well. The expected outcome is depicted by the dashed line in the right-side panel of Figure 2. The variance of p Bernoulli trials is $pq(1-q)$. Adjusted to this setting, we have for all $p \in \mathbb{Z}$:

$$var(\psi(p)) = 4|p|q(1-q). \tag{4}$$

If two points in the mapping are known, a bridge between them forms. Beliefs over the outcome for each policy on the bridge are determined by the relative number of paths between the ends of the bridge that pass through that outcome. The variance is determined similarly.

To demonstrate the main insight of the paper, we offer the simplest possible case. Suppose that $\psi(0) = -\gamma$ and let the politicians possess ideal points with the minimum possible polarization; i.e., set $\gamma = 1$. With a continuous hold on power the Democrat will not experiment and implement policy 0 in both periods. Proposition 4 shows that when his power is ephemeral, he deviates from this program and experiments in the first period.

Proposition 4 *For $\gamma = 1$ in sequence $D-R$, $p_1^* = 1$ for $q \in \left(0.646, \frac{3}{4}\right)$. In period 2, $p_2^* = 1$ if $\psi(p_1^*) = 0$, and $p_2^* = 0$ if $\psi(p_1^*) = -2$.*[13]

[13] The analytic value for the bound is: $\frac{4-\sqrt{2}}{4} \approx 0.646$. We ignore non-generic cases in this and all subsequent results.

This result illustrates the types of considerations that drive strategic experimentation in the setting considered in the previous section. If D experiments and the policy fails (and delivers outcome -2), then R will get stuck and choose to backtrack to the policy 0 and D obtains his ideal outcome in the second period. Since preemptive experimentation can head off a large experiment from R in the second period, it is in D's interest to preemptively experiment in this setting.

Proposition 4 establishes that D preemptively experiments for issues of moderate complexity. For the most simple and most complex issues preemptive experimentation is not profitable. Behavior in these cases is as follows.

Proposition 4b For $\gamma = 1$ in sequence $D - R$, $p_1^* = 0$ *for* $q \notin \left(0.646, \frac{3}{4}\right)$. *The second period policy is:*

$$p_2^* = \begin{cases} 0 \\ 1 \\ 2 \end{cases} \text{for } q \in \begin{cases} \left(\frac{1}{2}, \frac{5}{8}\right) \\ \left(\frac{5}{8}, 0.646\right) \\ \left(\frac{3}{4}, 1\right] \end{cases}.$$

Proposition 4b is interesting for why preemptive experimentation fails to be optimal. The Democrat's optimal choice depends on the Republican's anticipated response in the second period. For the most complex issues D stays at the status quo as the issue is too noisy even for R to experiment. For more moderate issues, $q \in \left(\frac{5}{8}, 0.646\right)$, R only experiments a little in the second period and there is no benefit to D to preempt. For the high q of the simplest issues, in contrast, R experiments boldly whether a first period experiment succeeds or not. As the issue is sufficiently predictable, D is confident that any experiment by R will succeed, so he prefers to get his own ideal outcome in the first period and, most likely, R's ideal outcome in the second period.

The effect on policy choice of different sequences of power brings us back to the question of voter welfare: Does taking turns in power serve the interests of voters? Proposition 5 shows that the answer is yes: The voter's ideal policy program is implemented whenever taking turns induces a preemptive policy experiment.

Proposition 5 *For* $\gamma = 1$ *and* $q \in \left(0.646, \frac{3}{4}\right)$, *the voter's ideal policy program is implemented if and only if the sequence of power is* $D - R$.

The proof of the proposition requires equilibrium policy choice for sequences $R - R$ and $R - D$. These, as well as the voter's complete ideal policy program, follow from straightforward computation and are contained in the appendix. In sequence $R - R$ the Republican experiments more than the voter desires, and sequence $R - D$ induces inefficient experimentation as regardless of the first period outcome the Democrat reverts to policy 0 in the second period. This leads to the striking conclusion of Proposition 5 that for issues of moderate complexity, the preferences of the voter line up with neither candidate, yet by having them take turns in power the voter is able to induce her ideal policy program. Proposition 5 is also of note in that it stands in stark contrast with the earlier finding that the voter never obtains her ideal policy program when the policy mapping is known (Lemma 1). This contrast exposes an important trade-off: Incomplete knowledge of the mapping is costly in that it allows for unintended consequences from policy experiments, yet it is precisely the prospect of such mistakes that tempers the policy choices when power is temporary and does so in a way that serves the interests of the voter. Proposition 5 shows that the benefit of policy moderation outweighs the cost of uncertainty and unintended consequences, leaving the voter strictly better off when knowledge of the policy mapping is incomplete.

We now describe first period behavior for $\gamma = 2$ in Proposition 6.[14]

Proposition 6 *For $\gamma = 2$ in sequence $D - R$,*
(i)

$$\begin{array}{lll} p_1^* = 1 & \text{for} & q \in [0.565, 0.618] \text{ and } p_2^* = 0 \text{ if } \psi(1) = -3, \\ p_1^* = -1 & \text{for} & q \in [0.619, 0.664] \text{ and } p_2^* = -1 \text{ if } \psi(-1) = -1. \end{array}$$

(ii) $p_1^* \in \{1, 2\}$ and $p_2^* > p_1^*$ for all realizations of $\psi(p_1^*)$, for all

$$q \in [0.674, 0.693] \cup [0.722, 0.749] \cup [0.768, 0.822] \cup [0.895, 0.926].$$

(iii) $p_1^* = 0$ otherwise.

The first part of Proposition 6 illustrates the logic behind the fact that exploration may either to the right or the left of the status quo in this game for moderate levels of polarization. When $q \in [0.565, 0.618]$, R will get stuck if D experiments to the right and the experiment fails, so it is in D's interest to preemptively experiment in the first period. For values of $q \in [0.619, 0.664]$, R will no longer get stuck if D experiments to the right and the experiment fails, but R will get stuck if D experiments to the left and the experiment results in a surprising outcome. As a result it is now in D's interest to experiment to the left in an attempt to get R stuck. This illustrates why experiments to the left can take place for larger degrees of polarization.

The rich variety of preemptive experimentation continues as the polarization of the politicians increases further. For $\gamma = 3$ the Democrat preemptively experiments in the hope of getting his successor stuck for all $q \in [0.543, 0.604]$ and uses experiments -1, 1, -2, and 3. (A full statement of the equilibrium is in the appendix.)

We conclude this section by illustrating an example that shows the intuition for the type of equilibrium behavior considered in Proposition 3. When there is a centrist status quo, alternating power instead provides an incentive for the incumbent to experiment less aggressively than he would if power were held continuously. This is illustrated in the following proposition:

[14] The calculations for this proposition were performed with a matlab program at increments of q of size 0.001. The program is available as an online appendix.

Proposition 7 *For the discrete policy space set $\gamma = 1$ and $\psi(0) = 0$. With a one-period horizon, D experiments with the policy -1 for all $q \in (q', 1)$. With a two-period horizon and the sequence $D - R$, he experiments with the policy -1 for all $q \in (q'', 1)$ where $q'' > q'$.*

The intuition for this is as follows: If D held power continuously and D experimented in the first period, then D would be able to keep the results of the experiment for two periods if the experiment resulted in D's ideal policy outcome, but D can change to some other policy if the experiment moves the policy outcome to R's ideal point. By contrast, if power alternates, then D will have to deal with the results of the experiment for two periods if the experiment results in R's ideal point, but R will want to change policies if the experiment results in D's ideal policy. Since D will no longer be able to benefit from experiments that result in his ideal policy outcome for multiple periods but will suffer for multiple periods if the experiment results in R's ideal outcome, D is less willing to experiment when power alternates under a centrist status quo. Thus alternating power is again beneficial for a centrist voter when there is a centrist status quo.

5 Discussion

Empirical Motivation.

A natural question to ask is whether preemptive experimentation is an empirical phenomenon. Systematic evidence is obviously difficult to obtain as essential features of the model are unobservable (such as policy preferences). Nevertheless, anecdotal evidence is available and is encouragingly supportive.

Evidence from politics in the United States is confounded by the separation-of-powers institutional structure: Is the small change to policy a preemptive experiment or is it the outcome of bargaining between the President and Congress? This problem is particularly acute on domestic policy where the President's powers for unilateral action are limited. Nevertheless, the experience of welfare reform in the 1990's under President Clinton is arguably best-interpreted as a preemptive experiment. Since Reagan, public opinion had tended toward the conclusion that welfare policy was overly generous and too 'left wing.' Bill Clinton used his time in office to reform welfare in a relatively modest way and significantly less than Republicans at the time were demanding. This preemptive reform lead to a more centrist outcome – a policy success – and the absence of further movement on welfare policy post-Clinton is consistent with the model's prediction that a small success will preempt the need for any further rightward movement.

Foreign policy offers the U.S. president considerably more freedom to act and offers a better opportunity to confront the model with data. Perhaps the best examples, in this regard, involve a reinterpretation of classic events in diplomacy. The famous episode of Nixon going to China in the 1970's is often interpreted through a signaling equilibrium in which Nixon held superior information about the "state of the world" (Cukierman and Tommasi 1998). Nixon's visit can easily be reimagined as a preemptive experiment, one intended to moderate policy and reveal information to his successors. An alternative history would involve Jimmy Carter using his time in office to make the first move toward Chinese integration and it is easy to conceive that this would have taken a much different form. My model offers a way to reinterpret Nixon's policy move without the need to believe that Nixon privately held superior information about the outcome of China policy.[15]

Nixon going to China is only the most famous example of this type of unexpected policy choice and many other policy surprises, both in the U.S. and elsewhere, have been interpreted in this way (see Rodrik (1993) and Cukierman and Tommasi (1998) for numerous examples). Examples from outside the U.S. arguably provide the strictest tests of the theory as in countries with parliamentary systems the leader of government is afforded considerably more policy freedom.

One such episode from Israel – that hasn't, to my knowledge, been affixed the Nixon goes to China label – is instructive.[16] After championing settlements in the West Bank and Gaza throughout his long career, Ariel Sharon reversed course when Prime Minister in 2004-2005 and successfully orchestrated Israel's unilateral withdrawal from Gaza in late 2005. This was a significant enough movement toward the opposition parties that it caused a breakdown in the ruling government coalition. The subsequent Hamas takeover of Gaza and the recommencement of rocket attacks into Israel – evidence of a policy failure – served to harden Israeli public opinion toward the Palestinians and cement Israeli conviction to not compromise on the West Bank. This learning and change in opinion is evident in opinion polls at the time.[17] It is possible to interpret Sharon's evacuation of Gaza not as a policy desiderata *per se*, but rather as an opportunity to credibly teach successor governments of different ideological bent that the much more dramatic step of ceding the West Bank to the Palestinians will not produce a desirable outcome.

Application to Democratization.

Empirical evidence for getting stuck may also be sought in the application of the model to democratization. It is frequently lamented that democracy does not always stick when it is introduced, often backsliding into dictatorship. Acemoglu and Robinson (2000, 2001, 2006) suggest that these transitions are due to fluctuating power conflicts among elites and the masses, in combination with the commitment power of political institutions. They argue that partial democratizations can be seen as preemptive compromises by the elite to head off the greater threat of social unrest and revolution. In this account, poor economic performance provides the "policy window" that makes the threat of revolution credible and, conversely,

[15] These machinations suggest the possibility for strategic polarization as it is the difference between Nixon and his successor (Jimmy Carter as it turned out) that, according to this interpretation, drove Nixon's actions. The model suggests why it was in the interests of Democrats to pose as pro-China to induce a preemptive experiment from Nixon, yet at the same time it explains why a limit to this strategy exists as if the Democrats' were too polarized Nixon would not attempt to preempt them. This possibility offers an interesting twist on the impact of the polarization recently appearing – and often lamented – in U.S. politics (McCarty, Poole and Rosenthal 2006).

[16] I thank Moses Shayo for suggesting this example.

[17] See http://www.spirit.tau.ac.il/xeddexcms008/manage.asp?siteID=5&lang=1&pageID=2342&stateID=65.

the threat of coup credible after democracy is established.

The informational model analyzed here offers an alternative interpretation of the history of democratization that complements existing accounts. The premise of this informational account is that the superiority of democracy for the masses is not common knowledge and is something that can only be learned by trial-and-error. This uncertainty and the need for experiential learning resonate with the historical experience. Uncertainty about democracy's effectiveness is evident ex-ante in Plato's argument that rule by an undemocratic "Philosopher King" is the ideal form of government. And the lessons of experience are evident ex-post in Churchill's (1947) famous quip:

> *"Many forms of Government have been tried and will be tried in this world of sin and woe. No one pretends that democracy is perfect or all-wise. Indeed, it has been said that democracy is the worst form of government except all those other forms that have been tried from time to time."*
> (Official Report, House of Commons (5th Series), 11 November 1947, vol. 444, cc. 206–07.)

More substantially, Ober (2008) provides a detailed account of the different political systems in use in ancient Greece and argues that the preeminence of Athens among the city-states was due to its singular reliance on democracy.[18] Viewed through this lens, the process of democratization can be thought of as a policy experiment (a large and significant experiment, to be sure, but an experiment nonetheless). As such, the retention of democracy is assured in the same way as for other experiments – by success.

The model then provides a framework to think through the process of democratization and strategic information revelation. A literal application of the model omits many first order details, yet still resonates with several prominent features in the history of democratization. Consider a dictator facing the certain prospect of being deposed by a revolutionary movement (with, obviously, different preferences). Further suppose that the dictator is currently enjoying the fruits of dictatorship (i.e., his ideal outcome). This corresponds to the situation captured by sequence $D - R$ and $\psi(0) = -\gamma$ that is the workhorse of the above analysis. Consequently, just as a Democrat preemptively experiments to shape the Republican's actions, a dictator will preemptively partially democratize to avoid the larger experiment of full democratization once he is deposed. Viewing partial democratization as strategic preemption is the same line of argument proposed by Acemoglu and Robinson, although here the channel for how preemption influences future behavior is exclusively informational.

Moreover, the informational theory provides an explanation for what comes after initial democratization. The correlation of economic underperformance and the failure of democracy can be interpreted as the natural reaction to a failed experiment.[19] Democratization

[18]The superiority of democracy does not seem to be a settled matter, however, as there is no shortage of popular commentators (and some academics) who argue that State Capitalism – as practiced most prominently in China – will dominate in the coming century.

[19]The frequent transitions in-and-out of democracy, such as experienced in Argentina, is not consistent

will only "stick" if it works, and if it doesn't work – the experiment is a failure – the populace possibly gets "stuck" in its hunt for a better political system and backslides into dictatorship. This backsliding may even be voluntary, explaining the puzzling phenomenon of dictators who overthrow democracy and assume power doing so with popular support.[20] On the other hand, if the democratization experiment is a success, democracy sticks and the populace updates positively about its effectiveness. This naturally leads to further democratization as a country moves further along the dimension of political freedom, as seen in the progressive deepening and evolution of democracy in the U.K. and many parts of western Europe.

This informational account is obviously a partial rendering and to the extent that information and learning about democracy drive democratization, they surely works in concert with other features of political conflict, such as power struggles, class conflicts, and institutional constraints. Nevertheless, the ability of information to speak to several prominent facts about the history of democratization, as well as the non-obviousness of democracy's allure across the sweep of history, suggest that experimentation and learning play some sort of role in the process of democratization.

Model and Results.
The model developed here is highly stylized and offers many opportunities for elaboration. The most obvious extension is to a longer horizon. The immediate challenge is a characterization of behavior over longer horizons. The Markov property of the random walk offers optimism that a value may be assignable to each information set and recursive methods applied. With a longer horizon, a preemptive experiment has more time to have an impact and the incentive to preemptively experiment may manifest *a fortiori*. A longer horizon also opens up the question of how exactly the politicians should optimally take turns in power. In practice, voters face issues of various and evolving complexity and the optimal sequence of political power over longer horizons may be irregular, following the vicissitudes of politics. This may involve eras of frequent transitions interspersed with lengthier periods of continuous power by one party, as has been observed in the modern history of the U.K. and several other countries.

The results reported in the previous section are built upon an asymmetry: That the known point in the mapping is more useful to D than to R. Although this is a natural state of affairs, it is possible that the status quo produces a more centrist outcome. In an online appendix I analyze the alternative case where the status quo outcome lies between D's and R's ideal points. Two interesting conclusions emerge from this case: First, that taking turns in power again impacts behavior relative to when power is continuously held, and, second, that the effect on behavior is reversed. With the status quo between the ideal outcomes of the players, any information revealed that is valuable to one player is ignored by the other as

with the deterministic outcomes in my model. Allowing observations of points in the mapping to be stochastic can generate (*a la* a Gittins index) the observed fluctuations in and out of democracy.

[20]A prominent recent example is Pervez Musharraf's rise to power in Pakistan. Famous historical episodes include Weimar Germany and inter-war Japan.

their interests are opposed. In contrast, any information revealed that is unfavorable to one player is picked up and used by the other player. Unusually, this generates a *negative* value for experimentation – the opposite of the standard intuition from active experimentation – and this negative value dampens the incentive of the early office holder to experiment with policy. Ironically, however, less experimentation is exactly what the median voter desires in this case, and having the parties take turns in office once again benefits the voter.

Applications Beyond Politics.
Worthy of independent exploration is the application of taking turns to decision making in organizations. For a boss managing a firm, the results suggest that he may only need to make the final decision on a project for his influence to cast a shadow over all previous choices of his workers, thereby easing his monitoring problem. An open question is then whether the boss' influence is enhanced or decreased as the horizon increases and where in the sequence the boss should optimally act. The results also suggest the boss is better off in some circumstances by designing himself out of the decision making process altogether, assigning workers with opposing biases to take turns in control, with the boss fitting into the role of the median voter.

Conclusion.
This note is well short of a fully developed theory of the strategy of information and policy dynamics. Nevertheless, it is a first step on a hard but important problem. Our intention is to shine light on a path that can be pursued. The possibilities are numerous and challenging. We have demonstrated in the simplest form how information can affect the incentives of decision makers when control is impermanent. Thus, in looking for how the past shapes the future, we need to look beyond institutions, budget deficits, and other formal constraints, to the informational environment that policies create and explore how information shapes the path of future choices.

References

[1] Acemoglu, Daron, Mikhail Golosov, and Aleh Tsyvinski. 2011. "Power Fluctuations and Political Economy." *Journal of Economic Theory* 146 (May): 1009-1041.

[2] Acemoglu, Daron, and James A. Robinson. 2000. "Why Did the West Extend the Franchise? Democracy, Inequality, and Growth in Historical Perspective." *Quarterly Journal of Economics* CXV: 1167-99.

[3] Acemoglu, Daron, and James A. Robinson. 2001. "A Theory of Political Transitions." *American Economic Review* 91 (September): 938-963.

[4] Acemoglu, Daron, and James A. Robinson. 2006. *Economic Origins of Dictatorship and Democracy*. Cambridge University Press (Cambridge, U.K.).

[5] Albrecht, James, Axel Anderson, and Susan Vroman. 2010. "Search by Committee" *Journal of Economic Theory* 145 (July): 1386-1407.

[6] Alonso, Ricardo. 2008. "Shared Control and Strategic Communication." Working paper.

[7] Alonso, Ricardo, Wouter Dessein, and Niko Matouschek. 2008. "When Does Coordination Require Centralization?" *American Economic Review* 98 (March): 145-179.

[8] Berry, Christopher R., Barry C. Burden, and William G. Howell. 2010. "After Enactment: The Lives and Deaths of Federal Programs." *American Journal of Political Science* 54 (January): 1-17.

[9] Callander, Steven. 2011. "Searching and Learning by Trial and Error." *American Economic Review* 101 (October): 2277-2308.

[10] Callander, Steven. 2011b. "Searching For Good Policies." *American Political Science Review* 105 (November): 643-662.

[11] Cukierman, Alex, and Mariano Tommasi. 1998. "When Does it Take a Nixon to Go to China?" *American Economic Review* 88 (March): 180-197.

[12] Dixit, Avinash, Gene M. Grossman, and Faruk Gul. 2000. "The Dynamics of Political Compromise." *Journal of Political Economy* 108 (June): 531-568.

[13] Gibbons, Robert. 2010. "Inside Organizations: Pricing, Politics, and Path Dependence." *Annual Review of Economics* 2: 337-65.

[14] Lindblom, Charles. 1959. "The Science of 'Muddling Through'." *Public Administration Review* 19: 79-88.

[15] McCarty, Nolan, Keith M. Poole and Howard Rosenthal. 2006. *Polarized America: The Dance of Ideology and Unequal Riches.* MIT Press.

[16] Merton, Robert K. 1936. "The Unanticipated Consequences of Purposive Social Action." *American Sociological Review,* 1 (December): 894-904.

[17] Messner, Matthias and Mattias K. Polbron. 2012. "The Option to Wait in Collective Decisions and Optimal Majority Rules." *Journal of Public Economics* 96 (June): 524-540.

[18] Moe, Terry M. 1989. "The Politics of Bureaucratic Structure." In Can the Government Govern? (John E. Chubb and Paul E. Peterson, eds.) Brookings Institution (Washington, D.C.).

[19] Ober, Josiah. 2008. *Democracy and Knowledge: Innovation and Learning in Classical Athens.* Princeton University Press (Princeton NJ).

[20] Page, Scott E. 2006. "Path Dependence." *Quarterly Journal of Political Science* 1 (1): 87-115.

[21] Persson, Torsten, and Lars E. O. Svensson. 1989. "Why a Stubborn Conservative would Run a Deficit: Policy with Time-Inconsistent Preferences." *The Quarterly Journal of Economics* 104 (May): 325-345.

[22] Pierson, Paul. 2000. "Increasing Returns, Path Dependence, and the Study of Politics." *American Political Science Review* 94 (June): 251-267.

[23] Pierson, Paul. 2004. *Politics in Time: History, Institutions, and Social Analysis.* Princeton University Press (Princeton NJ).

[24] Rodrik, Dani. 1993. "The Positive Economics of Policy Reform." *American Economic Review Papers & Proceedings* 83 (May): 356-361.

[25] Strulovici, Bruno. 2010. "Learning While Voting: Determinants of Collective Experimentation." *Econometrica* 78 (May): 933-971.

[26] Tabellini, Guido, and Alberto Alesina. 1990. "Voting on the Budget Deficit." *American Economic Review* 80 (March): 37-49.