

Agenda Control Under Policy Uncertainty*

Steven Callander[†]

Nolan McCarty[‡]

August 23, 2022

Abstract

Models of agenda-setting are central to the analysis of political institutions. Elaborations of the classical agenda setting model of Romer and Rosenthal (1978) have long been used to make predictions about policy outcomes and the distribution of influence among political actors. While the canonical model is based on complete and perfect information about preferences and policy outcomes, some extensions relax these assumptions to include uncertainty about preferences and reversion points. We consider a different type of uncertainty: incomplete knowledge of the mapping between policies and outcomes. In characterizing the optimal agenda setting under this form of uncertainty, we show that it amends substantively the implications of the Romer-Rosenthal model. We then extend the model dynamically and show that rich dynamics emerge under policy uncertainty. Over a longer horizon, we find that agenda control suppresses the incentive of legislators to experiment with policy, leading to less policy learning and worse outcomes than are socially efficient.

*For helpful comments we thank Keith Krehbiel, Salvatore Nunnari, Jon Bendor, seminar audiences at Stanford GSB and the Midwest Political Science Association meetings, and the Editors and referees of this journal. Ashutosh Thakur provided excellent research assistance.

[†]Graduate School of Business, Stanford University, Knight Management Center, Stanford, CA 94305; sjc@stanford.edu.

[‡]School of Public and International Affairs, Princeton University, Princeton, NJ 08540; nmc-carty@princeton.edu.

1 Introduction

A fundamental insight of positive political theory is that agenda power matters. She who decides *what* alternatives may be considered has as much power, if not more, than those who actually choose among those alternatives. This idea was first formalized in the classic analysis of school budgets in Oregon of Romer and Rosenthal (1978). The structure of the model developed in that paper provided the foundation for a vast literature that uses formal analysis to understand the strategy and logic of policymaking.

The Romer-Rosenthal model has come to be known as the Agenda Setter Model, or, more simply, the Setter Model. Its structure is parsimonious. Given a *status quo policy*, one player – the *Proposer* – has the opportunity to suggest an alternative policy to be considered. Another player, the *Voter*, may accept or reject, but not amend that proposal. If she accepts the proposal, it becomes the new policy. If she rejects it, the status quo remains in effect.¹ Thus, essentially the Setter Model is a take-it-or-leave-it game between the Proposer and the Voter that is played for a single round.

The elegance in Romer and Rosenthal’s result is that its simple structure vividly demonstrates the power of the Proposer over the final outcome despite the Proposer holding no formal voting power.² Beyond the conclusion that possession of *agenda power* is valuable, Romer and Rosenthal’s model reveals several predictions about the exact nature of this power and how it translates into policy choices and outcomes. These basic predictions are the building blocks of a vast formal literature on policymaking.

The Romer-Rosenthal model, as originally conceived and generally applied, supposes a world of certainty. The Proposer and the Voter both know the set of policies available to them and understand perfectly the outcomes each policy produces and how those outcomes affect their welfare. The assumption of complete certainty is a useful and sensible modeling tool yet it is clearly unrealistic. In practice, policymakers face considerable uncertainty about the outcomes that are produced whenever they change policy.

Our objective is to examine more closely the predictions of Romer and Rosenthal while relaxing their assumptions regarding information about the policy environment. Our results further reinforce their insight that agenda power is valuable. Yet we suggest that the connections between proposal power and policy choices, outcomes, and the welfare of the players, are richer and substantively changed by policy uncertainty.

Specifically, we suppose that the legislators know the full set of policies available to them but they are unsure as to which outcome each policy produces. To capture this uncertainty,

¹We abstract from the set of all voters and private market and taxation decisions that enrich Romer and Rosenthal’s account but that are not necessary for understanding the primary political mechanism.

²Of course, the model might allow the Proposer to vote as part of a larger legislature or electorate. The point remains that holding the agenda power magnifies his influence over the final outcome.

we use a recent model of policy uncertainty that was introduced in a series of papers by Callander (2008, 2011). This approach uses the Brownian motion stochastic process to capture policymaking uncertainty. We describe this model in detail in the following sections.

Adding uncertainty to the Agenda Setter Model is not itself novel. The analysis of uncertainty and its role in economic and political behavior has been the central enterprise of academic research for the past few decades. Where we depart from the literature is in focusing on the difficulty of policymaking faced by all legislators rather than asymmetric uncertainty that creates a signaling game between the players.³

In a world of policy certainty, the classic Romer-Rosenthal model delivers three key insights that follow from the power of agenda control.

- #1. Policy change is possible if and only if the status quo outcome is outside the ideal points of the legislators.
- #2. The Proposer does weakly better, and often strictly better, when the status quo is more extreme.
- #3. The Voter does weakly worse, and often strictly worse, when the status quo is more extreme, although his downside utility is limited.

The first property represents the foundation for the classic “gridlock interval.” Policy change happens only when both legislators agree on which direction to move. The second and third properties reflect the power of agenda control. The more extreme the status quo, the more leverage the Proposer has, and the better the policy outcome is for her. As this leverage is over the Voter, the worse outcome does not benefit the Voter and, in fact, makes him worse off as the more extreme status quo weakens his negotiating position. These properties are intuitive and ingrained throughout the literature on legislative policymaking. We show, however, that all three properties require amendment and are occasionally upended when we allow for policy uncertainty.

Policy uncertainty changes these predictions because, as is the case in practice, outcomes are no longer realized with certainty. This uncertainty is costly to legislators and tempers their preferences. As a result, legislators are not willing to change policy just because the status quo is not perfect. They are only willing to change if they can obtain something substantially better so that the benefit outweighs the risk. This requirement leads to more

³Such as the famous model of Gilligan and Krehbiel (1987) in which the Proposer has better information about the implications of policy choices than the Voter, or the models with uncertainty about Voter preferences such as Cameron (2000), Groseclose and McCarty (2001), and Rosenthal and Zame (2022). One of the few models that incorporate symmetric uncertainty is Buisseret and Bernhardt (2017) in which legislators have symmetric uncertainty about the preferences of future Proposers and Voters. This dynamic uncertainty does not affect the predictions of the one period model as does policy uncertainty.

inertia in policymaking and, therefore, a gridlock interval that is wider than the distance between the legislators' ideal points over outcomes, thereby changing Property #1 above.

This inertia constrains the ability of the Proposer to exploit her agenda power. Not only is she able to change policy less often, but when she does she moves it less than in the classic setting of policy certainty. That is to say, her *leverage* is decreased by policy uncertainty. In fact, the Proposer's own distaste for risk implies that she never moves the expected policy outcome all the way to her own ideal point, even if unconstrained by the voter.

The decreased leverage of the Proposer under policy uncertainty changes the logic of agenda power, altering the insights from Romer-Rosenthal's classic model. With less leverage in many situations, the Proposer benefits less from a more extreme status quo. Moreover, because a more extreme status quo requires a bigger policy change, and because uncertainty increases in the novelty of a new policy, an increase in leverage brings with it an unavoidable cost of risk. We show that the benefits of leverage are limited and, in fact, quickly overwhelmed by the costs of uncertainty. For moderate and extreme status quos, both the Proposer and the Voter are made worse off as the status quo is made more extreme. Even when leverage offers some benefit, it only does so for moderately extreme status quos. For more extreme status quos, the utility of both legislators declines in leverage and does so without bound. This substantively alters Properties #2 and #3 from the classic model.

Despite the canonical status of the Romer-Rosenthal model, there have been few efforts to directly test its predictions about optimal proposals and outcomes. An important exception is Clinton (2012) on Congressional voting over changes to the minimum wage laws.⁴ Clinton reports two main findings. First, he finds that policy change occurs far less often than would be predicted by the gridlock intervals computed from legislative preference estimates. Second, Clinton finds that when policy changes do occur, they are generally much smaller than those predicted by any of the agenda setting models. Specifically, the new policy outcomes do not move all the way into the gridlock interval defined by the ideal points of the pivots. These findings align directly with those of our model. To our knowledge, our model is the only one in the literature that makes the joint prediction of expanded gridlock intervals and smaller policy changes.⁵

The inclusion of uncertainty about policy leads naturally to the question of whether the

⁴Changes to the minimum wage laws fit the assumptions of our models well. For example, economists commonly estimate the employment effects of minimum wage increases from changes in state-level minimum wage laws. But from 1974 to 2021, the median real minimum wage increase for such changes is only 3.8%. (Author calculations based on Vaghul and Zipperer (2021)). By contrast, recent calls to raise the federal minimum wage to \$15 per hour would involve more than doubling the wage in the 20 states for which the current federal wage binds. Clearly, such a change involves far more uncertainty than those smaller revisions that form the basis of current knowledge. See Neumark and Shirley (2021) for a recent discussion of the uncertainties surrounding the employment effects of the minimum wage.

⁵For example, a model with costly proposal making would predict an expanded gridlock interval, but would not predict smaller policy changes conditional on a successful proposal. See the discussion in Section 3.2.5.

legislators can reduce uncertainty and the extent of their incentives to do so. The key insight is that the costs of policy uncertainty in the model fall most heavily on the Proposer and, thus, she has a greater incentive to reduce it. This observation bears an important implication for institutional design. It implies that if the legislators differ in their ability to reduce uncertainty, the one who is most able to reduce uncertainty should be assigned the role of proposer.

This observation corresponds to the arguments of Howell and Moe (2016, 2020) about needed reforms to the US separation of powers system. To Howell and Moe, a central liability of the US system is that agenda control is formally vested in the US Congress despite the fact that it is the executive branch that commands the greater capacities in policy planning and formulation. In their view, policymaking would be vastly improved if the executive branch formulated proposals on which the Congress accepted or rejected on an up-or-down vote. Our model provides some formal justification of this argument.

We then extend the model to a second period of policymaking. This captures any incentives that the Proposer and Voter might have to experiment with different policy choices. The second period opens up the opportunity for the players to experiment by bargaining over a first period policy, observe its consequences, and then bargain over a revision to that policy in a second period. The predictions of the two-period Brownian model stand in direct contrast with repeated versions of the Romer-Rosenthal complete information model (e.g. Primo 2002) or dynamic models of veto bargaining based on incomplete information about preferences (Cameron, 2000; McCarty, 1997). We find that agenda control sharply reduces the incentives for the agents to experiment with policy. The suppression of policy experimentation reflects that the agents may disagree about what constitutes a failed experiment. A good outcome for one player may be bad for the other. When such disagreement occurs, experimentation stops. Consequently, when the players have divergent preferences, the preferences of the multi-period game converge to those of the one-shot game.

2 Agenda Power Under Policy Uncertainty

We consider the classic Setter Model of Romer and Rosenthal (1978) amended only to include policy uncertainty. After introducing the Setter Model, we describe our model of policy uncertainty.

2.1 The Setter Model

The classic Setter Model is between two legislators: a Proposer (P) and a Voter (V). Both legislators care about outcomes, and they choose policies that are translated into outcomes by the mapping, ψ . The policy space and the outcome space are both given by the real line,

such that $\psi : \mathbb{R} \rightarrow \mathbb{R}$, and a policy p produces outcome $\psi(p)$.

Throughout we assume that the Proposer’s ideal outcome is $s < 0$ and the Voter’s ideal outcome is 0. For simplicity, we assume preferences are represented by a quadratic loss function in the outcome space. That is, utility for the Proposer and the Voter, respectively, are given by:

$$\begin{aligned} u(p) &= -(s - \psi(p))^2 \\ v(p) &= -\psi(p)^2. \end{aligned}$$

The policy-outcome space is depicted in Figure 1 and the ideal outcomes of the legislators are marked. A status quo policy, p_0 , is in place at the beginning of play, which produces outcome $\psi(p_0) > 0$.

Timing:

The timing of the Setter Model is simple. Within a legislative period:

1. The Proposer offers a policy $p_1 \in \mathbb{R}$.
2. If the Voter agrees (votes “yes”), policy p_1 is implemented and the outcome is $\psi(p_1)$, otherwise policy p_0 remains in place and the outcome is $\psi(p_0)$.

To put it another way, the Proposer designs the menu $\{p_0, p_1\}$ and the Voter selects a policy from the menu to implement. If the Voter is indifferent, we suppose that he selects the Proposer’s policy p_1 . (To retain policy p_0 the Proposer need only offer p_0 .)

The Setter Model is attractive in its simplicity. In the classic model, the policy mapping ψ is known and both legislators operate with complete information. Our focus is on agenda control when this assumption does not hold.

2.2 Modeling Policy Uncertainty

To capture policy uncertainty, we represent the policy mapping as the realized path of a Brownian motion. This follows the approach in Callander (2008, 2011) in that we assume that the legislators know the drift and the variance of the Brownian motion, μ and σ^2 , respectively, though they do not know the realized path. The legislators also know the outcome of the status quo policy, $\psi(p_0)$. We assume the drift is negative, $\mu < 0$, and the variance is, by necessity, positive, $\sigma^2 > 0$. Both players hold the same information such that neither has an informational advantage. One possible policy mapping is depicted in Figure 1 for $p_0 = 0$.

The Brownian motion representation concisely captures the richness of policymaking in practice and allows us to calibrate the model to the degree of complexity on any particular issue. It is also surprisingly tractable.

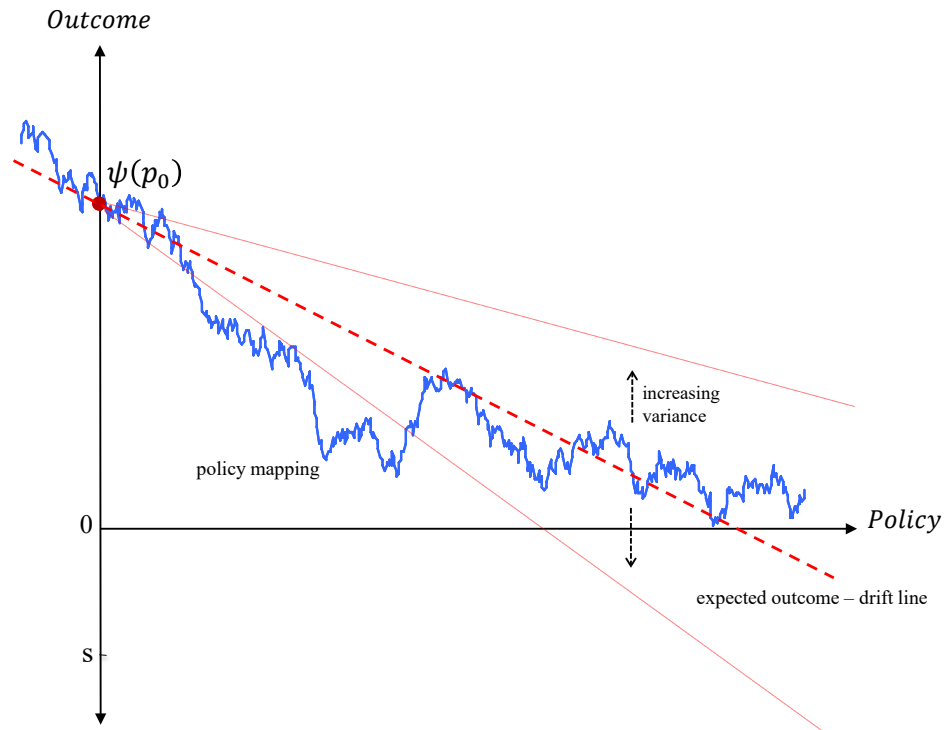


Figure 1: Brownian motion policy mapping with the drift line.

Expected outcomes under Brownian motion are shaped by the drift and variance parameters—the *theoretical knowledge* (Callander, 2011) held by the legislators—combined with practical knowledge of the status quo point—their *factual knowledge*. The expected outcome of all policies are given by the drift line of slope μ that is anchored at the status quo point. This is depicted as the dashed line in Figure 1. For a policy $p \in \mathbb{R}$, the expected outcome is:

$$E(\psi(p)) = \psi(p_0) + \mu p. \quad (1)$$

The drift parameter μ measures the expected rate of change. Thus, knowledge of the drift informs the legislators about which direction to move in order to shift the outcome in a particular direction and rate in expectation.

This is only in expectation, however, and for all policies other than the status quo, the outcome is unknown until a policy is tried and the outcome is observed. Beliefs over possible outcomes are normally distributed, with mean given by the drift line, and variance increasing in the distance a policy is from the status quo. This captures the idea that uncertainty is

increasing in the distance from what is known.⁶ Formally, the variance for a policy $p \in \mathbb{R}$ is:

$$\text{var}(\psi(p)) = |p| \sigma^2. \quad (2)$$

The variance measures the noisiness of the policy mapping. The higher the variance, the less predictable are policy outcomes. The ratio of variance to drift measure the uncertainty that must be accepted for each unit shift in the expected policy outcome. We define half of this value as the *complexity* of the policy issue, denoted by α , such that $\alpha = \frac{\sigma^2}{2|\mu|}$.

The combination of quadratic utility and normally distributed outcomes deliver a concise mean-variance representation for utility. For any policy experiment, expected utility can be written as utility at the mean of the distribution less the variance:

$$Eu(p) = -[s - E(\psi(p))]^2 - \text{var}(\psi(p)).$$

An appeal of the Brownian motion representation is that it captures several important features of politics that are otherwise obscured. One important feature is that of unintended consequences. Ever since Merton's (1936) famous promulgation of this into a law, the idea that policies intended to produce one outcome may produce a different outcome, and indeed may shift the outcome in the opposite of the intended direction, has been widespread. The Brownian motion captures this possibility. As is evident in Figure 1, the policy mapping is not monotonic, often shifting directions as policy changes. It is not entirely unpredictable, however, as the outcome is more likely to shift in the intended direction—the direction of the drift—than not.

Alternative Approaches to Policy Uncertainty

Modeling policy uncertainty is, of course, not new. The most well-known approach to modeling uncertainty is that of Gilligan and Krehbiel (1987) from their study of legislative expertise where they introduce the now-familiar formulation

$$x = p + \omega,$$

where ω is distributed uniformly over $[0, 1]$. Importantly, this is fundamentally the same structure that we employ. The difference is in the simple versus complex issues that each approach captures. In Gilligan and Krehbiel (1987), the policy mapping is linear and with a known slope of one. The legislators lack knowledge only of the intercept term. Critically,

⁶It is important to note that uncertainty is relative to what is known rather than what is the status quo. For simplicity, we assume these are the one point. This distinction matters for the relationship between this model of policy uncertainty and a model of direct costly policy change. We discuss this connection explicitly in Section 3.2.5.

if the legislators knew the status quo point, they would know the entire mapping. Thus, to have a known status quo point and legislative uncertainty about the policy mapping, a richer structure is required.⁷

The classic model of Romer and Rosenthal (1978) implicitly employs this same structure. The model assumes a known status quo point and that both legislators possess full knowledge of the policy mapping. Importantly, full knowledge of the mapping implies full control over policy outcomes. The legislators can accurately and precisely obtain any desired outcome by simply identifying the policy to which that outcome corresponds and implementing it. Our model with policy uncertainty relaxes this precision, which is the source of differences in the equilibria between our model and those in Romer and Rosenthal (1978).

3 Agenda Power in a Single Period

3.1 The Certainty Benchmark

Lemma 1 states the classic Romer-Rosenthal result. For this statement, we assume that the policy mapping has slope -1 and is given by: $\psi(p) = \psi(p_0) - p$. With perfect knowledge of the policy mapping, the Proposer exploits her agenda control to move the outcome as close to her ideal as possible without losing the support of the Voter.

Lemma 1 (*Romer and Rosenthal, 1978*) *Suppose the policy mapping is known by both players and given by $\psi(p) = \psi(p_0) - p$. In equilibrium, the Proposer offers p_1^* , where:*

(i) *For $\psi(p_0) > 0$, $p_1^* = \min\{2\psi(p_0), \psi(p_0) - s\}$, such that $E\psi(p_1^*) = \max\{-\psi(0), s\}$, and the Voter accepts.*

(ii) *For $\psi(p_0) \in [s, 0]$, policy does not change.⁸*

(iii) *For $\psi(p_0) < s$, $p_1^* = \psi(p_0) - s$, such that $E\psi(p_1^*) = s$, and the Voter accepts.*

This simple lemma generates two fundamental insights that have evolved into conventional wisdom about bargaining over policy outcomes. The first regards policy stability. Lemma 1 provides the foundation for the *gridlock interval* that has become central to our understanding of legislative politics, particularly that of the U.S. (Krehbiel, 1998). The gridlock interval

⁷This is not a problem for Gilligan and Krehbiel (1987) as they do not assume a known status quo and focus on advice by an expert committee (who knows the mapping) and an uniformed legislative body. However, uncertainty in their model lasts only a single period because the true mapping is revealed with the outcome of the first-period policy choice. With the Brownian motion structure, the true mapping is never fully revealed, even after an arbitrary number of policy choices. See Callander (2011) and the two period extension in Section 4.

⁸The equilibrium is not unique in this case although all equilibria are outcome equivalent.

describes the location of policy outcomes that are entrenched. In our model, this is the policies in the interval $[s, 0]$ between the legislators' ideal outcomes. In the U.S. context, the gridlock interval is taken as the boundary between the most extreme of the House median, the President, and the Senate filibuster and veto override pivots. Woon and Cook (2015) uses this measure to make quantitative predictions about policy change in the 111th Congress with the election of Barack Obama to the presidency.

The second insight is about policy change. The logic of case (i) is Romer-Rosenthal's famous "flipping" strategy. The Proposer offers the policy that flips the outcome across the Voter's ideal point leaving her indifferent, up to the point at which the Proposer obtains her own ideal outcome. The Proposer uses her power to make a take-it-or-leave-it offer to extract all of the surplus created by reforming policy. The Voter benefits from the new policy as well only when the Proposer is able to obtain her ideal point and there are no further gains from leverage.⁹

This classic logic is depicted in Figure 2. The policy mapping is the straight blue line of slope -1.¹⁰ In the figure, the Proposer offers policy p_1^* that produces outcome $\psi(p_1^*)$ with certainty, leaving the voter indifferent between it and policy p_0 .

We focus hereafter on the case of $\psi(p_0) > 0$. In this case, the insights of Romer-Rosenthal yield three lessons. The first is the simplest: The Proposer has leverage and policy change occurs whenever $\psi(p_0) > 0$.

Property #1: For every $\psi(p_0) > 0$, policy changes from the status quo: $p_1^* > 0$ and the Voter accepts.

The second lesson is that the more extreme the status quo outcome, the more power the Proposer has and the higher her equilibrium utility.

Property #2: The Proposer's utility is weakly increasing in $\psi(p_0)$ for $\psi(p_0) > 0$. Specifically, $\frac{du}{d\psi(p_0)} > 0$ for $\psi(p_0) \in (0, -s)$ and $u = 0$ for all $\psi(p_0) \geq -s$.

For status quo outcomes only a little beyond the Voter's preference, the Proposer has only limited leverage, and the more extreme the status quo, the more leverage she has which makes her strictly better off. For a sufficiently extreme status quo, however, the Proposer's leverage is sufficient to obtain her ideal outcome. At that point, she is no better or worse off if the status quo is more extreme.

⁹The interests of the legislators are aligned in case (iii) and both benefit from policy change.

¹⁰The Romer-Rosenthal intuition is typically depicted purely in outcome space. The policy choice and the underlying policy mapping are left implicit as they are trivial. We depict them here so as to facilitate the extension to uncertain policy environments and the Brownian motion representation.

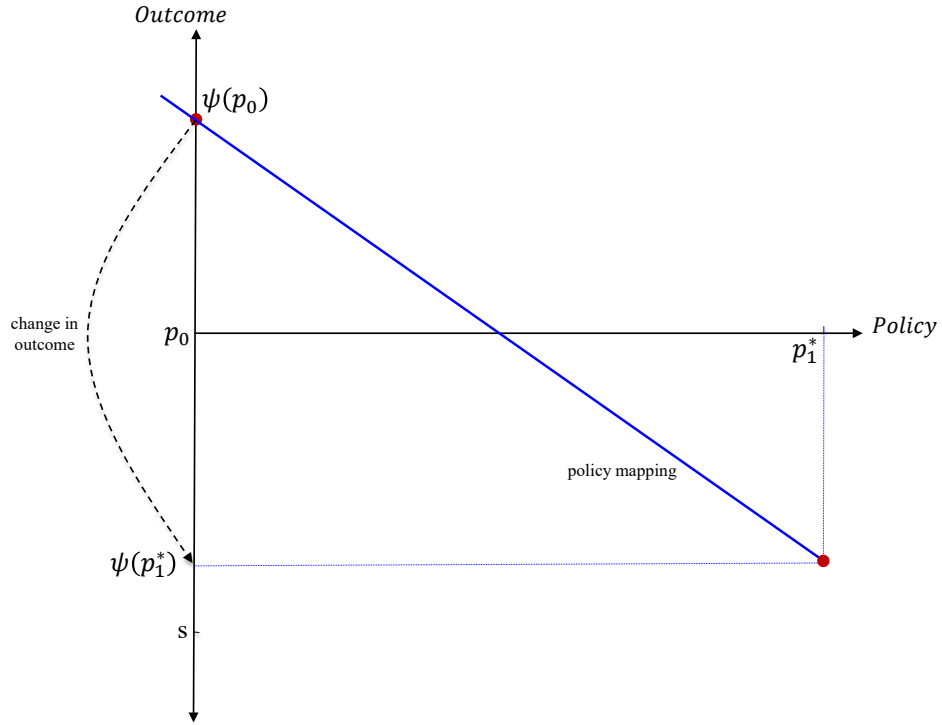


Figure 2: Equilibrium Policy Under Certainty (Romer-Rosenthal).

The third lesson is that the converse holds for the Voter. He is worse off under more extreme the status quo outcomes, and his utility strictly decreases until the Proposer obtains her ideal outcome.

Property #3: The Voter's utility is weakly decreasing in $\psi(p_0)$ for $\psi(p_0) > 0$. Specifically, $\frac{dv}{d\psi(p_0)} < 0$ for $\psi(p_0) \in (0, -s)$ and $v = -s^2$ for all $\psi(p_0) \geq -s$.

As the status quo outcome moves further from the Voter's ideal, the Voter's utility from the status quo declines. Because the Proposer is able to extract all of the benefit from the policy change, the Voter's ex post utility is also worse under the more extreme the status quo. This holds up until the Proposer obtains her ideal outcome, at which point the Voter's utility is constant at $-s^2$ regardless of the status quo.

3.2 The Setter Model Under Uncertainty

How is the equilibrium of Lemma 1 — and the lessons drawn from it — affected by the presence of uncertainty? We now return to the Brownian motion specification and proceed in steps through the derivation of equilibrium. We focus hereafter on the case where $\psi(p_0) > 0$ and show that the three Properties identified above change significantly in the presence of policy uncertainty.

3.2.1 Voter Indifference

A key property of Romer-Rosenthal that continues to hold with policy uncertainty is that the Voter is left indifferent by the Proposer unless the Proposer is able to obtain approval of her most preferred policy choice. We begin by characterizing the policy proposal that achieves Voter indifference for each possible status quo outcome. Recall that we define policy complexity as $\alpha = \frac{\sigma^2}{2|\mu|}$.

Lemma 2 *For status quo policy p_0 and outcome $\psi(p_0)$, the Voter is indifferent over policies p^I and p_0 where $p^I = \frac{2}{|\mu|}(\psi(p_0) - \alpha)$ if $\psi(p_0) > \alpha$. For $\psi(p_0) \leq \alpha$, the Voter strictly prefers the status quo policy over all other policies.*

Proof. Using the mean-variance representation for utility,

$$Ev(p) = -[\psi(p_0) + \mu p]^2 - |p|\sigma^2 \quad \text{and} \quad v(p_0) = -\psi(p_0)^2.$$

Clearly, $v(p_0) > Ev(p)$ for any $p < 0$. For $p > 0$, $Ev(p) \geq v(p_0)$ if and only if

$$p \leq -\frac{2\psi(p_0)\mu + \sigma^2}{\mu^2} = \frac{2}{|\mu|}(\psi(p_0) - \alpha)$$

Note that the set of p preferred to the status quo is empty if $\psi(p_0) < \alpha$. This is because the cost of the uncertainty of changing policy ($|p|\sigma^2$) dominates the benefit of a better expected outcome ($-2\psi(p_0)\mu p - \mu^2 p^2$) for all $p > 0$ so that the Voter strictly prefers the status quo over all other policies. For $\psi(p_0) \geq \alpha$, the Voter is indifferent between the status quo and $p^I = \frac{2}{|\mu|}(\psi(p_0) - \alpha)$. ■

Two features of Lemma 2 are worth emphasizing. The first is that the Voter strictly prefers the status quo — and depriving the Proposer of leverage — not only at the Voter's ideal outcome, but for a range of outcomes beyond his ideal. Specifically, for a policy with outcomes between 0 and α , there is no policy proposal that the Proposer can induce the Voter to support.

The second important feature of Lemma 2 is the location of the policy proposal that leaves the voter indifferent. Recall from Lemma 1 that, under policy certainty, the indifferent policy for the Voter is that with outcome exactly symmetric to the status quo outcome around the Voter's ideal outcome. Specifically, for status quo outcome $\psi(p_0)$ and Voter ideal of 0, the indifferent policy has outcome $-\psi(p_0)$. The indifferent policy flips across the Voter's ideal outcome. This is not the case with policy uncertainty. The indifferent policy flips expected outcomes not over 0, but across outcome $\alpha > 0$. (Recall that in Lemma 1 it is assumed that $\mu = -1$.)

These two features are the consequence of the Voter’s response to risk. In the range of $0 < \psi(p_0) \leq \alpha$, there are obviously outcomes that the Voter prefers (namely his ideal outcome), yet any policy that may be selected to change the outcome entails risk. The outcome may move in the intended direction, though it may move too far, or it may move in the unintended direction *a la* Merton (1936). These concerns, combined with his risk aversion, cause the Voter to consider outcomes within α of his ideal to be preferred over any risky policy. Those readers familiar with Callander (2011) will recognize this critical value as the threshold for a “good enough” outcome in the search for good policies.

3.2.2 Equilibrium Proposals

The Voter’s policy preferences under uncertainty are interesting in their own right, but our main interest is on their implications for the Proposer’s strategy and the equilibrium policy outcomes. The key insight here is that the Voter’s willingness to tolerate imperfect outcomes and his aversion to engaging in risk constrains the Proposer and limits her leverage. Proposition 1 describes equilibrium behavior.

Proposition 1 *Under policy uncertainty, the equilibrium policy, p_1^* , is given by:*

- (i) $\psi(p_0) \leq \alpha$. No policy change occurs and the status quo remains in place.¹¹
- (ii) $\psi(p_0) \in (\alpha, -s + \alpha)$. $p_1^* = p^I$, the Voter accepts, and $E\psi(p_1^*) = -\psi(p_0) + 2\alpha$.
- (iii) $\psi(p_0) \geq -s + \alpha$. $p_1^* = \frac{1}{\mu}(s - \psi(p_0) + \alpha)$, the Voter accepts, and $E\psi(p_1^*) = s + \alpha$.

Proof. Case (i) is drawn immediately from Lemma 2. If the Voter strictly prefers the status quo, he cannot be induced to vote for any alternative policy.

In cases (ii) and (iii) the Proposer uses her leverage against the Voter, and the policy changes. The optimal behavior and distinction between the cases follows from two facts: (a) When changing policy is optimal, the Proposer’s utility increases in p until reaching a maximum when the expected outcome is a distance α short of her ideal point so that $E\psi(p_1^*) = s + \alpha$ (see Callander (2011) for the details). (b) As expected utility is concave in policy, Lemma 2 implies that the Voter agrees to any policy in the interval (p_0, p^I) .

Combining these facts, the Proposer offers policy p^I unless her optimal is less than this. The two cases follow by identifying the p such that $E\psi(p) = s + \alpha$. ■

Case (i) reflects the Voter’s willingness to tolerate imperfect outcomes. There may be outcomes that both he and the Proposer prefer, but they cannot implement them with certainty and the risk involved in trying exceeds the benefit. This feature implies that gridlock is more pervasive under uncertainty. The gridlock interval extends beyond the ideal points of the

¹¹There are multiple equilibria in this case although all are outcome equivalent.

pivotal legislator, the measure used throughout the literature theoretically (Krehbiel, 1998) and empirically (Woon and Cook, 2015). Thus, Property 1 from Romer-Rosenthal changes fundamentally, and we amend it as follows.

Property #1^u: For $\psi(p_0) \in [0, \alpha]$, $p_1^* = 0$ and gridlock holds. Only for $\psi(p_0) > \alpha$, does $p_1^* > 0$, the Voter accepts, and policy changes from the status quo.

For complex issues, the range from 0 to α can be large such that policy entrenchment is pervasive, and policy uncertainty increases gridlock substantially.

In cases (ii) and (iii) the Proposer does have leverage and is able to shift (expected) policy in the direction she prefers. Unlike in classic Romer-Rosenthal, however, the Proposer is not able to flip the outcome across the Voter’s ideal. Rather, in case (ii), she is able to flip the expected outcome only across the threshold, α . Thus, risk not only limits *when* the Proposer has leverage, it limits the *strength* of her leverage when she has it.

The status quo is sufficiently extreme in case (iii) that the Proposer is not constrained by the Voter’s preference and is able to obtain her ideal policy choice. Nevertheless, she does not set the expected outcome of the policy to her ideal outcome. This decision parallels the Voter’s policy preferences in Lemma 2. The Proposer shifts policy to the point at which the expected outcome is α above her ideal, accepting the imperfect expected outcome as a suitable trade-off against the additional risk that would be required to shift the expected outcome closer to her ideal. This is the “good enough” outcome for the Proposer.¹²

Figure 3 depicts the equilibrium expected policy outcome as a function of $\psi(p_0)$, the status quo outcome. The red line is the classic Romer-Rosenthal solution under certainty, as stated in Lemma 1. The black line is the equilibrium with policy uncertainty from Proposition 1. It is evident that the Proposer’s leverage under uncertainty holds only for more extreme status quo outcomes and is weaker where it does exist. Moreover, the figure shows the extended gridlock region (the first segment of the black line) where the Proposer and the Voter both prefer a set of policy outcomes to the status quo, yet are deterred from changing policy by uncertainty.

3.2.3 Proposer Utility

Proposition 1 reports equilibrium policies and expected outcomes but does not report the equilibrium utility of the legislators. In the classic Romer-Rosenthal environment with certainty, there is a tight link between voter indifference, the Proposer’s leverage, and utility. These connections change substantially when there is policy uncertainty.

To explore the link between leverage and utility in the presence of policy uncertainty, we begin with the Proposer’s utility as the status quo outcome becomes more extreme. With

¹²That the same α holds for both legislators is due to the properties of quadratic utility and is not general.

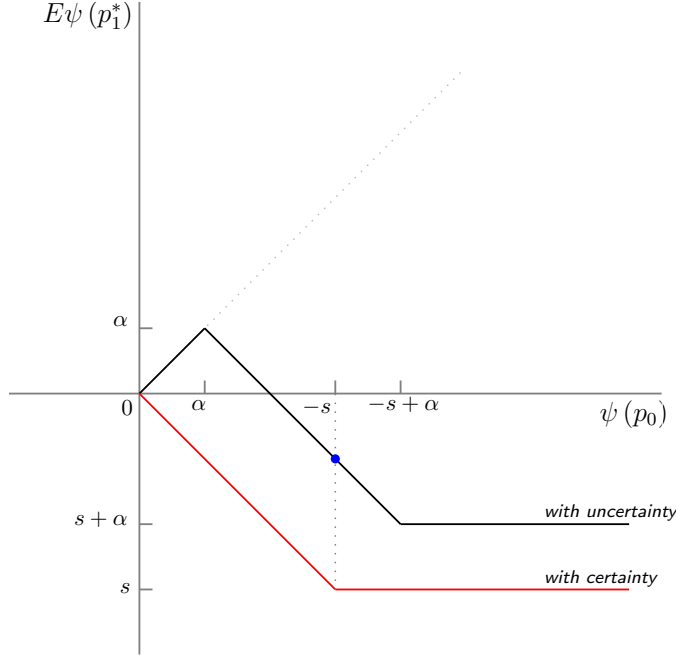


Figure 3: Equilibrium Expected Outcomes

policy certainty, a more extreme status quo outcome often makes the Proposer strictly better off and never worse off. Under policy uncertainty, the relationship depends on the extremity of the Proposer's ideal outcome. The following corollary breaks down the three possible cases.

Corollary 1 *In the equilibrium of Proposition 1, for:*

- (i) $s \in (-\alpha, 0)$, the Proposer's expected utility is strictly decreasing in $\psi(p_0)$.
- (ii) $s \in (-4\alpha, -\alpha)$, the Proposer's expected utility is non-monotonic in $\psi(p_0)$ and maximized at $\psi(p_0) = 0$.
- (iii) For $s \leq -4\alpha$, the Proposer's expected utility is non-monotonic in $\psi(p_0)$ and maximized at $\psi(p_0) = -s$.

Proof. We begin with several observations from Proposition 1. Case (iii) implies that for $\psi(p_0) > -s + \alpha$, the Proposer's utility is strictly decreasing in $\psi(p_0)$. This relationship occurs because the expected outcome remains constant at $s + \alpha$ whereas the attendant variance is increasing in $\psi(p_0)$ which lowers the Proposer's utility. Case (i) implies that for $\psi(p_0) \leq \alpha$ the Proposer's utility is strictly decreasing in $\psi(p_0)$ as no policy change is possible and $\psi(p_0)$ is greater than the Proposer's ideal outcome s .

Case (ii) of Proposition 1 in which the Proposer has positive but imperfect leverage remains. The Proposer's utility at the optimum p_1^* is:

$$Eu(p_1^*) = -[s - 2\alpha + \psi(p_0)]^2 - 4\alpha[\psi(p_0) - \alpha].$$

Differentiating:

$$\begin{aligned}\frac{dEu(p_1^*)}{d\psi(p_0)} &= -2[s - 2\alpha + \psi(p_0)] - 4\alpha \\ &= -2(s + \psi(p_0)) \\ \frac{d^2Eu(p_1^*)}{d\psi(p_0)^2} &= -2 < 0.\end{aligned}$$

As this case requires $\psi(p_0) > \alpha$, if $s \in (-\alpha, 0)$, the first derivative is always negative, concluding the proof for case (i) of the corollary.

For $s < -\alpha$, the Proposer's utility is strictly increasing for $\psi(p_0)$ in the neighborhood of α , establishing the non-monotonicity claims of cases (ii) and (iii). The Proposer's expected utility is maximized when:

$$\frac{dEu(p_1^*)}{d\psi(p_0)} = 0 \Rightarrow \psi(p_0) = -s.$$

This delivers the Proposer a utility of:

$$\begin{aligned}Eu(p_1^*) &= -[s - (2\alpha + s)]^2 - 4\alpha[s - \alpha] \\ &= 4s\alpha.\end{aligned}$$

Comparing this utility to that for $\psi(p_0) = 0$, which is simply $-s^2$, establishes that utility is maximized at $\psi(p_0) = 0$ when $s \in (-4\alpha, -\alpha)$ and at $\psi(p_0) = -s$ when $s \leq -4\alpha$. ■

We know that the Proposer gains leverage over the Voter in all three cases whenever the status quo allows for policy change (i.e., $\psi(p_0) > \alpha$). Nevertheless, in cases (i) and (ii), the Proposer is strictly better off if the status quo is less extreme and her leverage to change policy is more circumscribed. Indeed, she is better off when the status quo outcome is at the Voter's ideal point—that is, inside and at the far end of the classic gridlock interval—and she has no leverage at all over the Voter. When the Proposer's preferences are moderate and not too dissimilar to the Voter's, leverage from an extreme status quo is a cursed sword. The Proposer uses and benefits from leverage when she has it, but the cost of that leverage is greater uncertainty. For moderate preferences, the risk outweighs the benefit and the Proposer is worse off.

Figure 4 depicts the Proposer's expected utility in equilibrium. The top and middle panels depict cases (i) and (ii), respectively. The red line is Proposer utility in the classic Romer-Rosenthal model with policy certainty. The Proposer's utility is negative with a status quo outcome of 0 and, under certainty, she gains utility monotonically as she gains leverage. Her utility increases in the status quo outcome until she can obtain her ideal outcome. The black

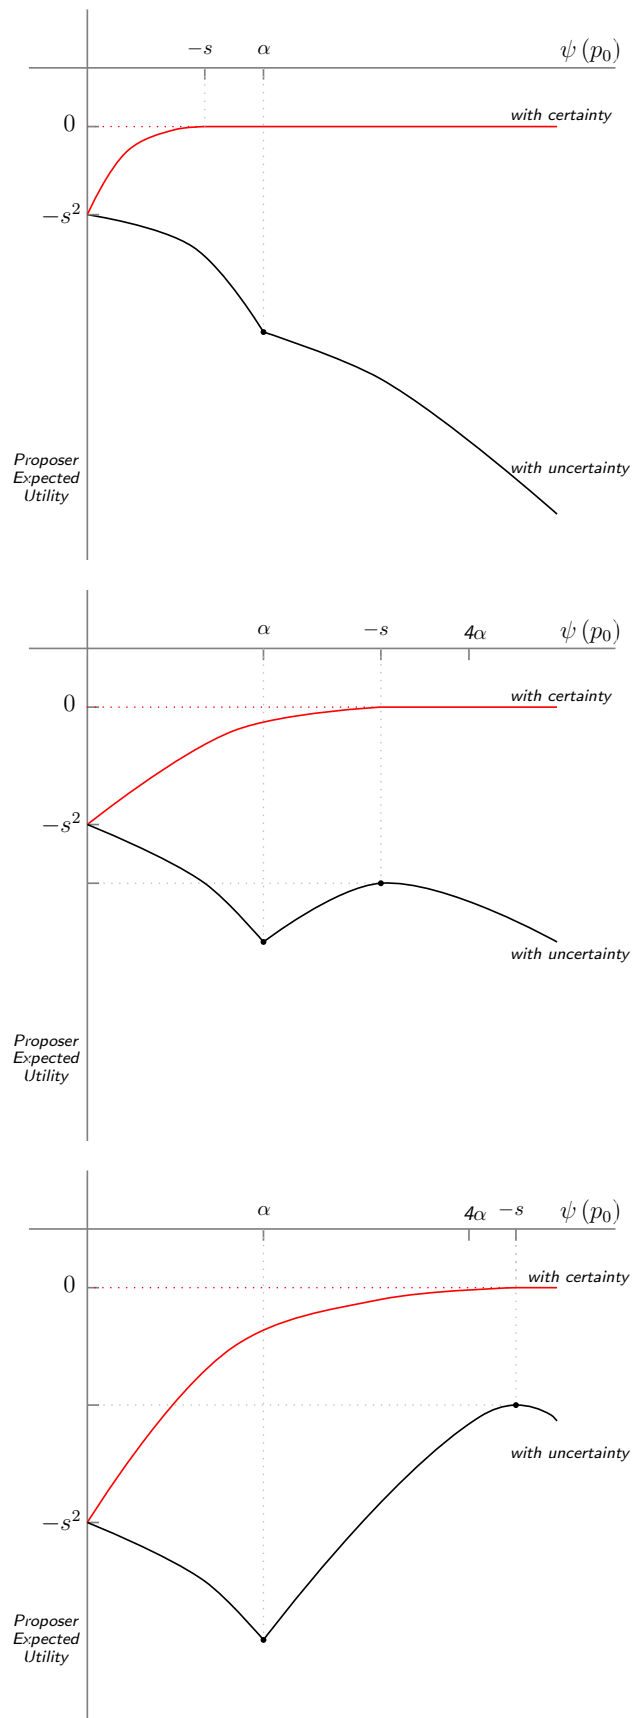


Figure 4: Equilibrium Expected Utility for the Proposer

curves depict Proposer utility under uncertainty. In the top panel of case (i), the Proposer's utility decreases monotonically without a lower bound. The kink point in utility is where the Proposer gains leverage over the Voter. In both cases (i) and (ii) the Proposer benefits from that leverage relative to what she would get from the status quo, but not so much to overcome the negative effects of the worse status quo.

Case (iii) is depicted in the bottom panel of Figure 4. In this case, the Proposer sufficiently benefits from leverage to overcome the negative effects of a worse status quo and can be better off with leverage than without.¹³ Intriguingly, the Proposer's utility is maximized when the status quo outcome is the exact mirror image of her ideal outcome (at $-s$), which is where she first obtains her maximum utility under policy certainty. This peak does *not* correspond to where her leverage is maximized, however. Thus, even when leverage is valuable, it is best in a limited quantity.¹⁴

Thus, even in case (iii) when the Proposer benefits from leverage, that benefit holds only for a narrow region of status quo outcomes. For very extreme status quos, the Proposer is strictly worse off than if the status quo provided no leverage and her loss is unbounded. This property, which holds in cases (i) and (ii) as well, contradicts that which arises under certainty. We amend Property #2 accordingly.

Property #2^u: With policy uncertainty, the Proposer's utility is strictly decreasing in $\psi(p_0)$ for all s , except for $s < -\alpha$ in the interval $\psi(p_0) \in [\alpha, -s]$.

Decomposing Agenda Power Under Uncertainty

Conceptually, there are two distinct channels through which policy uncertainty affects the Proposer's equilibrium utility. First, policy uncertainty may change the value of agenda control. Second, policy uncertainty imposes a cost for moving the status quo.

We can isolate the first channel by calculating the value of proposal power while keeping policy uncertainty constant. Specifically, we compare the Proposer's equilibrium utility with the utility she would obtain if the Voter set policy unilaterally under policy uncertainty. The cases in Corollary 2 follow those in Proposition 1, and the proof is relegated to the appendix.

¹³As the status quo outcome becomes more extreme, the Proposer's utility from it decreases quadratically. Her utility from the equilibrium proposal p_1^* decreases linearly. The value $s = -4\alpha$ is the crossover point where the linear decrease dominates the quadratic decrease.

¹⁴At $\psi(p_0) = -s$, $E\psi(p_1^*) = \alpha - (-s - \alpha) = 2\alpha + s$. The Proposer's leverage is maximized when $E\psi(p_1) = s + \alpha$, but $(2\alpha + s) - (s + \alpha) = \alpha > 0$.

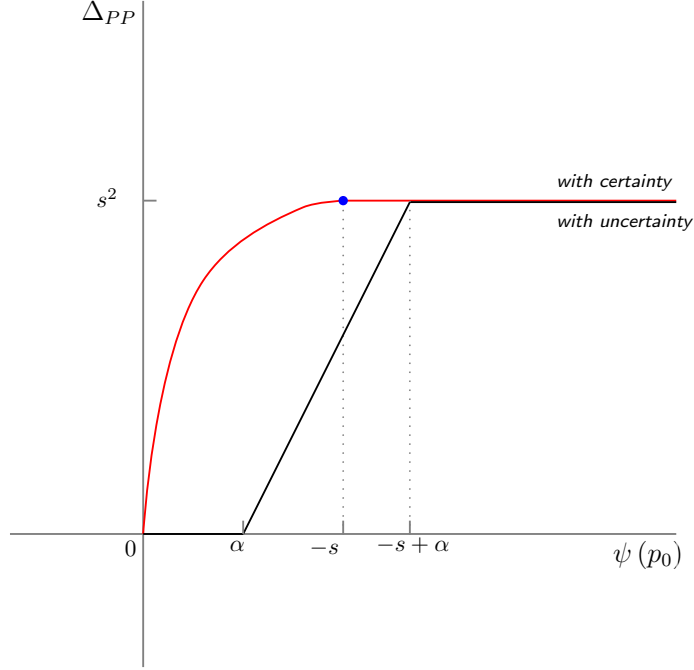


Figure 5: The Value of Proposal Power

Corollary 2 *The value of proposal power is given by Δ_{PP} , the difference in equilibrium utility and the counterfactual where the Voter sets policy. For:*

$$\begin{array}{ll}
 (i) & \psi(p_0) \leq \alpha, \\
 (ii) & \psi(p_0) \in (\alpha, -s + \alpha), \\
 (iii) & \psi(p_0) \geq -s + \alpha,
 \end{array}
 \quad
 \Delta_{PP} = \begin{cases} 0 \\ 2|s|(\psi(p_0) - \alpha) \\ s^2 \end{cases}$$

The value of Δ_{PP} is depicted in Figure 5 by the black line. The red line shows the corresponding value under policy certainty. Proposer power is less valuable under policy uncertainty when the Proposer has moderate preferences, but is equally valuable when the Proposer has extreme preferences.

The logic of this result is clear for case (i). In this case gridlock holds only under policy uncertainty. Thus, with policy uncertainty, the Voter is unwilling to change policy whether he controls it or the Proposer does and Proposal power is useless. In contrast, under policy certainty, the Proposer is able to flip policy even for such moderate status quos, and she benefits by shifting policy from the Voter's ideal closer to her own ideal point.

Although the value of proposal power in case (iii) is the same whether policy is certain or not, the details differ across the two environments. With policy certainty, the Proposer flips policy all the way to her ideal point and receives her ideal outcome (and utility of 0) versus the Voter's ideal outcome (and utility of $-s^2$). Under policy uncertainty the flip is of the same size, although it is instead from α to $s + \alpha$. This policy change is more valuable to the Proposer given she has quadratic utility. However, in addition to the policy change,

the Proposer now faces additional risk. Due to properties of quadratic utility, the added risk exactly cancels out the added benefit of the change in expected outcome.

Perhaps the most interesting is case (ii) where the Proposer must accommodate the Voter and has positive but limited leverage. As the Proposer is more constrained by the Voter under policy uncertainty, combined with the added cost of uncertainty from changing policy, the value of proposal power is lower under policy uncertainty. However, because the expected outcome is further from the Proposer’s ideal, the marginal benefit of flipping a more extreme status quo is higher, and eventually the value of proposal power catches that under policy uncertainty as $\psi(0)$ transitions to case (iii).

The upshot of Corollary 2 is that while the value of proposal power remains weakly positive and increasing in $\psi(p_0)$, policy uncertainty reduces the proposer’s advantage for non-extreme status quo outcomes, and for extreme status quos it only matches that under policy certainty. This implies that for extreme status quos, the large gap in Proposer utility in Figure 4 between the two lines is due solely to the cost of uncertainty. For more moderate status quos, the divergence between the two lines is a combination of uncertainty and weakened proposal power.¹⁵

3.2.4 Voter Utility

The Voter’s utility is also affected by policy uncertainty. The Proposer continues to leverage the Voter as much as she can and drives her to indifference, but just as the Proposer is negatively affected by the uncertainty this creates, so too is the Voter.

Corollary 3 *The Voter’s expected utility is strictly decreasing in $\psi(p_0)$.*

Proof. Consider the three cases of Proposition 1 and note that the Voter’s utility from the status quo outcome is strictly decreasing in $\psi(p_0)$. Case (i) is obvious since $\psi(p_0) > 0$ by assumption. Case (ii) is immediate from the equilibrium condition that the Voter is indifferent over the status quo and the proposal. For case (iii), note that the expected outcome is independent of $\psi(p_0)$ where the size of p , and thus the variance, is strictly increasing in $\psi(p_0)$. ■

¹⁵A natural question is to ask about the role of risk aversion in both legislators preferences. The quadratic-loss form is intuitive and analytically convenient, yet it does yield special properties. Nevertheless, the classic underlying logic of flipping policy about the Voter’s ideal outcome does not rely on risk aversion. If the Voter were less risk averse than the Proposer, then the Proposer, by still driving the Voter to indifference, would gain even more leverage. Nevertheless, as long as the Proposer herself is risk averse over outcomes, the shape of Proposer utility in Figure 4 would remain as eventually the Proposer will be unconstrained by the Voter, and the increased risk that follows from a more extreme status quo would lower her utility. (Note that risk aversion is the natural consequence of an environment in which a legislator has an internal ideal point and outcomes have full support, as they do here. Even if the Proposer had linear utility curves, uncertainty that covers her ideal point would generate risk aversion over outcomes.)

The Voter is strictly worse off as status quo outcomes become more extreme. Thus, the limits on leverage that uncertainty creates do not benefit the Voter. Rather, the limits of the Proposer's leverage purely reflect the social cost of uncertainty.

The impact on Voter utility relative to the classic Romer-Rosenthal environment with policy certainty is less profound than for the Proposer, but nonetheless important. The impact on the Voter is no longer bounded below for extreme status quos. The Voter's utility not only declines in the status quo for moderate values, it declines for all values. We amend Property #3 as follows.

Property #3^u: With policy uncertainty, the Voter's utility strictly decreases in $\psi(p_0)$. Specifically, $\frac{dv}{d\psi(p_0)} < 0$ for all $\psi(p_0) > 0$.

Combining Corollaries 1 and 3, we see that in most cases both legislators are hurt by making the status quo outcome more extreme. Beyond a critical point, specifically $\psi(p_0) = -s$, a more extreme status quo unambiguously hurts all legislators even though it delivers more leverage to the Proposer. This suggests that the efficacy of a political system to correct policy outliers is not as simple and as painless as one would conclude from the classic Romer-Rosenthal setting under certainty.

3.2.5 Costly Policymaking, Increasing Uncertainty, and Functional Forms

As we have highlighted, uncertainty in our model makes policy change costly. Indeed, in the one period model, the cost of uncertainty is equivalent to a model with direct cost of changing policy where the fixed cost of change is zero and the marginal cost increases in the size of the policy change.

Such a combination of costs yields predictions consistent with Clinton's (2012) two findings about legislative behavior. Notably, other combinations will not. For instance, a positive fixed cost of policy change will generate a wider gridlock interval (Clinton's first finding) but not the smaller increments when policy does change (Clinton's second finding). With the fixed cost paid, the Proposer follows the prediction of Romer-Rosenthal under policy certainty. Similarly, assuming no fixed cost but a marginal cost that is initially zero generates Clinton's second finding but not his first. Our model provides a microfoundation for why costly policy change generates the behavior that Clinton documents.

The equivalence of our model and one with a direct cost of policy change holds only in this simple, one period setting. Moreover, even in the one period model, equivalence relies on the status quo being the known point in the mapping. In our model of policy uncertainty, variance increases from the known point, whereas in a costly policy change model the cost (presumably) increases from the status quo. This can matter even in a one period model if a second point is known. For instance, considering a tax increase from a status quo of 20%

to 25% may be informed by the outcome of a 40% tax rate in a neighboring country. In such a setting, the cost of uncertainty follows a non-linear and potentially non-monotonic path. This richness emerges naturally in a multi-period setting, as we explore in Section 4. Thus, policymaking in the face of policy uncertainty is fundamentally different from a setting with a direct cost of policy change.¹⁶

3.3 Reducing Complexity

The preceding section highlights the distinctive impacts of policy uncertainty and complexity on the welfare of the Proposer and Voter. The upshot is that while complexity reduces the ability of the Proposer to change policy and exploit her bargaining power, the Voter is somewhat insulated from the effects of uncertainty given her option to retain the status quo. This finding may provide some insight into issues of institutional design.

To be concrete, suppose players could exert costly effort to reduce the complexity of the policy issue, α . For example, we might imagine that either the Proposer or Voter might invest in higher quality staff whose experience might be helpful designing policies that can be implemented with less noise. However, instead of modelling such a game, the main insights are available simply by comparing the marginal benefits of complexity reduction for both players.¹⁷ As the following corollary shows, the marginal benefit of complexity reduction is always at least as great for the Proposer than for the Voter.¹⁸

Corollary 4 *When $\psi(p_0) \geq \alpha$ The marginal benefit of a reduction of α is strictly greater for the Proposer than for the Voter. If $\psi(p_0) < \alpha$, the marginal benefit is zero for both players.*

The upshot of the corollary is that if both players are equally efficient at reducing uncertainty, the Proposer will have the greater incentive to contribute to that end. Second, if the players vary in their capacities to reduce uncertainty, endowing the Proposer role with the power to reduce uncertainty, and assigning the more efficient player to be the Proposer, would minimize policy uncertainty.¹⁹ This observation relates to the arguments of Howell and Moe (2016, 2020) about needed reforms to the US separation of powers system. To Howell and Moe, a central liability of the US system is that agenda control is formally vested in the US

¹⁶Not all models of policy uncertainty generate the same uncertainty, of course. An appealing property of the Brownian motion is that learning is local. For instance, trying minimum wages of \$7.00 and \$7.10 will reveal a lot about other wages around that level but little about the likely outcome of a minimum wage of \$15. In the Gilligan and Krehbiel framework, in contrast, the assumption of linearity means that observing any two points fully reveals the mapping, even when both the intercept and the slope are unknown.

¹⁷One possible form for this extended game might be $\alpha^* = \frac{\alpha}{1+e_p+e_v}$ and the players pay k_p and k_v per unit of effort. But such a model would provide no additional insights beyond comparing marginal disutilities of α .

¹⁸The proof of this, and later results, can be found in the appendix.

¹⁹As we do not model citizens or voters, we make no claim as to whether the minimizing uncertainty corresponds to welfare maximization.

Congress despite the fact that it is the executive branch that commands the greater capacities in policy planning and formulation. In their view, policymaking could be vastly improved if the executive branch formulates proposals on which the Congress accepts or rejects on an up-or-down vote. Our model provides a formal justification of this argument.

4 Agenda Power Over Time

Although policymaking may be parsimoniously modeled as a one-shot game, in practice it is ongoing. Policies are implemented, adjusted, and reversed. Perhaps after many iterations, policy stabilizes. In this section, we extend the Setter Model to a simple dynamic setting, namely to two periods.

This is an unusual extension of the Setter Model, not because it is technically difficult, but rather because in the classic setting it is relatively boring.²⁰ With policy certainty, repeating the Setter Model over two, or any number of periods with the same players is boring as nothing of interest happens after the first period. Specifically, the first period plays out exactly as described in Lemma 1, and then policy remains stable thereafter.²¹ Thus, the back-and-forth of policy and rich dynamics that are observed in practice do not emerge as an equilibrium to the model.²²

Formally, it is straightforward to extend the model to dynamic policy making. We adopt the game and timing as described in Section 2 and repeat it a second time, assuming that the policy implemented in period 1 becomes the status quo in period 2. Both legislators are forward looking and, thus, care about the outcome produced in both periods; they discount the future at rate $\delta \in [0, 1]$.

Dynamic Policy Experimentation

With a multiple period horizon under policy uncertainty, the Setter Model is now a problem of optimal policy experimentation. The classic insight of experimentation is that agents are more willing to experiment the longer their time horizon (Gittins, 1979). Thus, an agent facing a two-period horizon experiments more than an agent with only a single period of choice. The intuition for this result is simple. After experimenting in the first period, the agent can discard failures and retain successes. Thus, the benefit of finding a success in the

²⁰One exception is Buisseret and Bernhardt (2017) who show how interesting dynamics emerge when there is uncertainty about the preferences of the future Proposer and Voter.

²¹See Primo (2002) for the details. A similar one-change dynamic is evident in dynamic models of veto bargaining based on incomplete information about preferences (Cameron, 2000; McCarty, 1997).

²²In the models of veto bargaining with incomplete information about preferences, there may be multiple rounds of bargaining prior to striking an agreement, but the sequence of offers is monotonic. That feature does not hold in our model.

first period is enhanced as it can be enjoyed for twice as long, whereas failures remain only as costly as they are in the single-period game.

Our model of policy experimentation differs in two ways from those in statistics and economics. First, both legislators must agree to any experiment with one possessing agenda control.²³ Second, the legislators choose the novelty—and, thus, the riskiness—of an experiment and not merely the intensity of experimentation or even whether to experiment or not. In this setting, the classic insight of experimentation translates into a legislator preferring a larger policy change in the first period than if she had only a single-period horizon.

We show here that one of these differences matters and one does not. It remains true that legislators wish to experiment more boldly with a longer horizon and this translates into a preference for more risk and for policies that depart further from the status quo. However, because of agenda control and the need for both legislators to agree, the incentive to experiment is tempered in the longer horizon. In fact, we find that in some situations policy experimentation is suppressed to the point that policy choice is indistinguishable from what occurs in the one-period environment.²⁴

Equilibrium in Two Periods

To demonstrate this result, we employ a combination of numerical and analytical results.²⁵ Figure 6 depicts the classic experimentation intuition for the parameter values $\sigma^2 = .15$, and $\mu = -.15$ (so that $\alpha = .5$). The figure depicts the optimal first period policy choice for each legislator if that legislator had full control of policy (i.e., does not require the other legislator’s agreement). The two lines represent for each legislator what they would choose with a one-period horizon and with a two-period horizon. The Voter’s ideal policy does not depend on the Proposer’s ideal point, whereas the Proposer’s ideal policy increases in how negative is

²³This feature is shared with a large literature on dynamic legislative bargaining, although these models are of complete information without experimentation; classic references include Baron (1996), Kalandrakis (2004), and Penn (2009). We discuss the literature on collective experimentation momentarily.

²⁴This result resonates with the bias against experimentation found in the literature on collective experimentation. In that literature the inefficiency is due to uncertainty over the distribution of gains and losses from reform that, over time, creates a shifting majority. This means a pivotal voter today may not be pivotal tomorrow, rendering her wary of experimentation (Strulovici, 2010; Gieczewski and Kosterina, 2020). Voting in our paper is by unanimity and this concern is not relevant. Rather, inefficiency emerges due to the endogeneity of the status quo. In agreeing to a change today, the Voter may lock herself in to an unattractive policy tomorrow because she lacks the necessary agenda control to reverse the change (Messner and Polborn (2012) assumes the irreversibility of a policy reform exogenously). This brings us closer to Anesi and Bowen (2021), particularly their results on collegial voting rules. Relative to that paper, we allow for a continuous policy space and a richer experimentation technology (rather than binary policy) and we show how the suppression of experimentation manifests as smaller, more incremental policy change.

²⁵The two-period model is intuitive yet analytically difficult. The second period policy choice falls into one of four types of responses—we refer to these below as scenarios—and taking expectations requires cutting the normal distribution (of first period outcomes) into four pieces. This is easy to state, and easy to calculate numerically, though difficult to manipulate analytically.

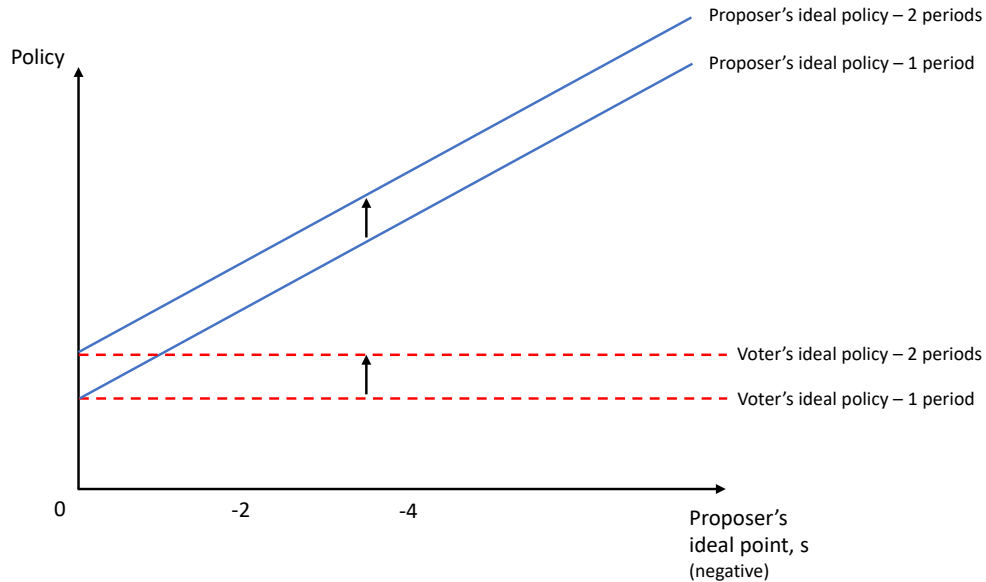


Figure 6: Optimal First Period Policy for Voter and Proposer Without Agenda Control

her ideal outcome. As can be seen, both legislators want to experiment more boldly with a longer horizon.

It would seem a reasonable conjecture that the equilibrium with agenda control would follow a similar shift. We show this is not the case. This is depicted in Figure 7. The lower solid red line depicts the one-period agenda control equilibrium from Proposition 1; recall, in equilibrium, policy shifts twice as much as the Voter's ideal shifts, up until the Proposer obtains her ideal point. The second solid red line performs the same shift, doubling the Voter's ideal policy for a two-period horizon. The actual equilibrium policy is, however, given by the yellow line. We denote equilibrium choices for the two-period horizon with a hat; thus, \hat{p}_1^* for the first period policy choice with a two-period horizon.

To understand this result, and see why agenda control suppresses experimentation, we return to the classic experimentation intuition that successes are kept and failures discarded. Consider then the moderate failure from the Voter's perspective in which the first period outcome overshoots the Voter's ideal outcome, ending up somewhere near the Proposer's ideal outcome.²⁶ If the Voter has full control, he will change policy again, reversing course, and aim to move the outcome back toward his ideal outcome. As the Proposer holds agenda power, however, this does not happen and the failure is not abandoned.

²⁶The legislators now know two points in the mapping. Between the known point a Brownian bridge forms and beliefs are an interpolation of the two points with variance increasing and then decreasing across the bridge. Outside of the known points, beliefs continue to follow Equations (1) and (2), albeit anchored at the nearest known point. See Callander (2011) for details.

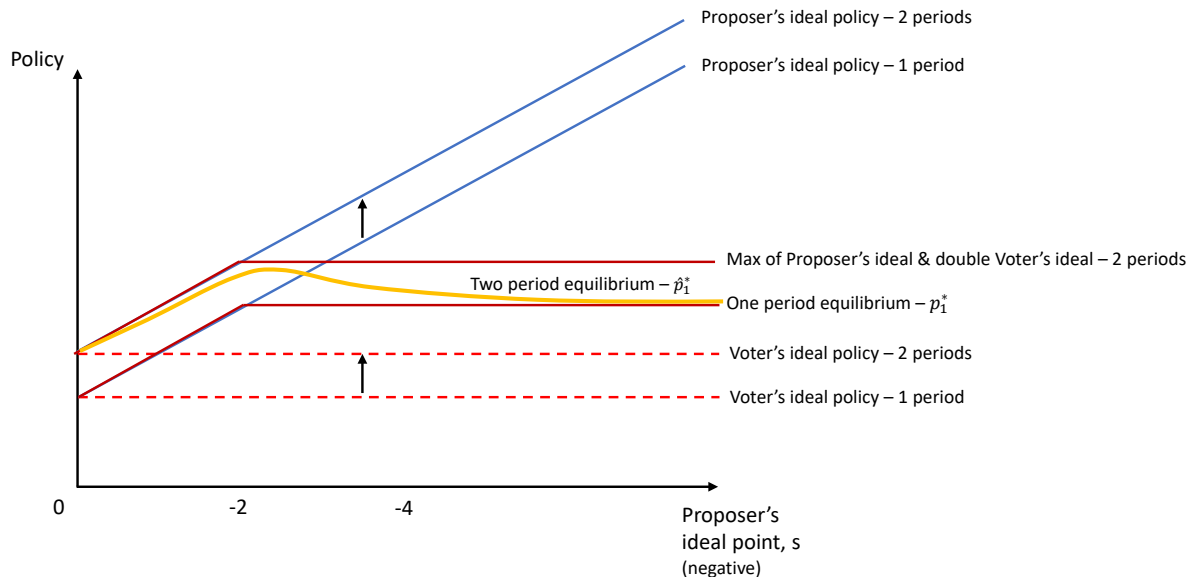


Figure 7: Two-Period Equilibrium with Agenda Control

A similar logic applies if the outcome shifts toward the Voter's ideal but not by enough. Here the Voter wants to change policy, pushing further in the same direction, in the hopes of getting closer to his ideal. The Proposer agrees, and offers to change policy, but she offers to change policy by much more than the Voter wants. In fact, as we saw in the one-period model, the Proposer will offer the Voter a policy that leaves him indifferent between it and retaining the first period policy. In this case, therefore, the failure is abandoned, but it is not replaced by something better, rather it is replaced by something exactly equivalent.

In each of these scenarios the Voter does not benefit from policy change in the second period. He receives a potentially different outcome in the second period, to be sure, but because it delivers the same utility as he receives in the first period. It is as if he had a one-period horizon. Figure 8 depicts these scenarios, marked as (3) and (2), respectively.

These are not the only possibilities, however. In the other two scenarios there is some common interest between the Voter and the Proposer, and the Voter benefits from policy change in the second period. If the first period outcome overshoots even the Proposer's ideal point, scenario (4) in Figure 8, the Proposer will reverse the direction of policy, and move it left such that the outcome moves back toward the center. She does not move policy as much as the Voter would like, but the change does benefit the Voter.

Scenario (1) represents the classic "law of unintended consequence" in which policy change that sought to move the outcome in one direction actually causes it to move in the opposite direction (Merton, 1936). With an outcome distant from both legislators' ideal points, there is agreement that, should a new policy be tried, it would be far to the right. Such a policy

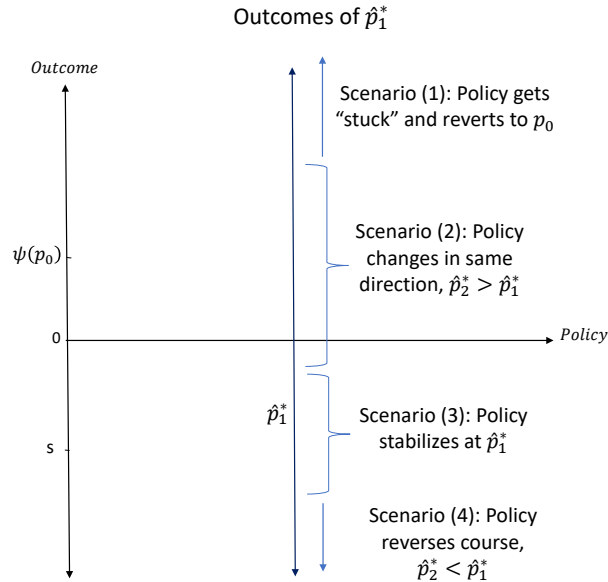


Figure 8: First Period Policy Outcomes & Second Period Policy Choice

would require large uncertainty, however, and this is costly to the legislators. As such, whilst such an experiment is more preferable than the first period policy, p_1^* , it is less preferable than policy p_0 . The initial status quo p_0 is not an attractive outcome, but it is without risk and, given the unintended consequence of the first period, it is the best choice. In the terminology of Callander (2011), learning gets “stuck” and the Proposer offers to reverse course and revert back to p_0 .

This may end learning but it is a relatively good outcome for the Voter. The failure of period 1 is abandoned, and he receives a policy that he strictly prefers to that failure, thereby benefiting from policy change in the second period. Consequently, this scenario gives him more encouragement to experiment in the first period.

Combining the four scenarios, the Voter’s willingness to experiment is suppressed from what he would choose on his own but not entirely. This result is evident in numerical simulations for a broad range of parameter values. It should be noted, however, that just as in the one-period model, this willingness to change policy does not benefit the Voter in equilibrium, as the Proposer uses her agenda power to offer a first-period policy that is favorable to her and drives the Voter to indifference.

A special case of this result emerges as the Proposer’s ideal outcome becomes increasingly extreme. As can be seen in Figure 7, the two-period equilibrium approaches the one-period equilibrium as the Proposer’s ideal outcome becomes very negative. We prove this result analytically.

Proposition 2 *As $s \rightarrow -\infty$, $\hat{p}_1^* \rightarrow p_1^*$.*

This result follows from the four scenarios described above. As the Proposer’s ideal point becomes more extreme, the gap in preferences between the two legislators becomes ever larger. Consequently, the two scenarios in which there is common interest between the legislators, (1) and (4), become very unlikely. With scenarios (2) and (3) dominating, the Voter does not benefit from policy change in the second period, and being unable to either discard failures or benefit from discarding them, his willingness to experiment is exactly as it is over a one-period horizon.²⁷

5 Conclusion

Policymaking in practice is difficult. Given the tremendous uncertainty about how the world works, the implications of a policy change can only be imperfectly predicted. This reality colors all efforts to strategize over policy choice as well as attempts to leverage other policymakers into favorable choices.

We have demonstrated in this paper how such policy uncertainty affects even basic intuitions as those that emerge from the simplest model of policy bargaining. The core insight is that leverage and policy payoffs are no longer tightly linked when there is uncertainty about the mapping of policy choices into outcomes. A more extreme status quo may well provide an agenda setter with more leverage, but also requires her to tolerate more risk in changing the policy. Such uncertainty can leave her worse off than if the status quo were moderate and she had no leverage at all. Our framework provides a richer understanding of agenda control and the tension between conflicting and common interests in legislative bargaining.

The properties of our model have rich implications for the practice of politics. We explored briefly the implications for incentives to reduce complexity and agenda control over a longer horizon. The striking result that experimentation and learning are suppressed when one player controls the agenda leads to questions of institutional structure. How have political institutions evolved, or how can they be designed, so as to ameliorate this inter-temporal inefficiency? More generally, Romer and Rosenthal’s timeless insight into the importance of agenda control provides the foundation for much work since across a variety of contexts. How policy uncertainty impacts the insights of these theories, say, for example, the legislative bargaining of Banks and Duggan (2000), is an important open question. Many avenues of

²⁷This logic is not dependent on the two-period horizon and extends to any length horizon. For less extreme Proposer ideal points, a longer horizon should facilitate more experimentation as the benefit of correcting bad outcomes increases in the horizon. It remains true, however, that whenever an outcome is realized within the classic gridlock interval of $[s, 0]$, experimentation ceases thereafter. A conjecture is that as the players’ patience increases to one, the gridlock interval of Property #1^u shrinks toward the classic gridlock interval. Establishing the truth of this conjecture, even numerically, is beyond our present capabilities. Even if true, the expansion of the gridlock interval in Property #1^u holds for any level of legislator impatience.

investigation present themselves, and incorporating policy uncertainty into our theorizing offers the promise of a deeper understanding of the policymaking process.

6 Appendix

Proof of Corollary 2 We can compute the utilities of the Proposer for the Voter-only game and compare to her expected utility for each of the three cases in Proposition 1. The Proposer's utility is

$$u^v = \begin{cases} -(\psi(p_0) - s)^2 & \text{if } \psi(p_0) < \alpha \\ -(s - \alpha)^2 - 2\alpha(\psi(p_0) - \alpha) & \text{if } \psi(p_0) > \alpha \end{cases}$$

The equilibrium utilities from the three cases of proposition 1.

$$u = \begin{cases} -(\psi(p_0) - s)^2 & \text{Case i} \\ -(s + \phi(p_0) - 2\alpha)^2 - 4\alpha(\psi(p_0) - \alpha) & \text{Case ii} \\ -\alpha^2 - 2\alpha(\psi(p_0) - \alpha - s) & \text{Case iii} \end{cases}$$

Let $\Delta_{pp} = u - u^v$ be the utility difference. Therefore,

$$\Delta_{PP} = \begin{cases} 0 & \text{Case i} \\ -2s(\psi(p_0) - \alpha) & \text{Case ii} \\ s^2 & \text{Case iii} \end{cases}$$

Cases (i) and (ii) are obvious. Canceling terms reduces case (iii) to $-\alpha^2 + 2s\alpha + (s - \alpha)^2 = s^2$.

Proof of Corollary 4: The claim can be verified by differentiating the utilities from the three cases in Proposition 1.

Case 1: $\psi(p_0) \leq \alpha$. The respective utilities for the Proposer and Voter are:

$$\begin{aligned} u(\psi(p_0); \alpha) &= -(s - \psi(p_0))^2 \\ v(\psi(p_0); \alpha) &= -\psi(p_0)^2 \end{aligned}$$

Clearly, $\frac{\partial u}{\partial \alpha} = \frac{\partial v}{\partial \alpha} = 0$. Therefore, neither player has an incentive to reduce complexity.

Case 2: $\psi(p_0) \in (\alpha, -s + \alpha)$ and $p_1^* = -\frac{2\psi(p_0) + \sigma^2}{|\mu|^2}$. Now:

$$\begin{aligned} u(\psi(p_0); \alpha) &= -(2\alpha - \psi(p_0) - s)^2 - 4\alpha(\psi(p_0) - \alpha) \\ v(\psi(p_0); \alpha) &= -(2\alpha - \psi(p_0))^2 - 4\alpha(\psi(p_0) - \alpha) \end{aligned}$$

Differentiating with respect to α , we obtain

$$\begin{aligned} \frac{\partial u}{\partial \alpha} &= 4s < 0 \\ \frac{\partial v}{\partial \alpha} &= 2(-2\alpha + \psi(p_0)) - 2\psi(p_0) + 4\alpha = 0. \end{aligned}$$

Since $s < 0$, $\frac{\partial u}{\partial \alpha} < \frac{\partial v}{\partial \alpha}$. Since both partials are non-positive, the Proposer has the larger benefit of complexity reduction.

Case 3: $\psi(p_0) \geq -s + \alpha$.

$$\begin{aligned} u(\psi(p_0); \alpha) &= -\alpha^2 - 2\alpha(\psi(p_0) + s - \alpha) \\ v(\psi(p_0); \alpha) &= -(s + \alpha)^2 - 2\alpha(\psi(p_0) + s - \alpha) \end{aligned}$$

Therefore,

$$\begin{aligned} \frac{\partial u}{\partial \alpha} &= 2\alpha - 2\psi(p_0) - 2s \leq 0 \\ \frac{\partial v}{\partial \alpha} &= 2\alpha - 2\psi(p_0) - 4s \end{aligned}$$

Since $s < 0$, $\frac{\partial u}{\partial \alpha} < \frac{\partial v}{\partial \alpha}$. Since both partials are non-positive, the Proposer has the larger benefit of complexity reduction.

Proof of Proposition 2. By generalizing Proposition 2 in Callander (2011) to allow for any s , the threshold that divides scenarios (1) and (3) is the solution to: $\psi(p_1) = \frac{1}{2\alpha} [\psi(p_0)^2 - 2s\psi(p_0) + \alpha^2]$. This is increasing without bound as $s \rightarrow -\infty$, and thus, for a fixed p , the probability of scenario (3) approaches 0. Similarly, a necessary condition for scenario (4) is that $\psi(p_1) < s$, and, for fixed p , the probability of scenario (4) goes to 0 as $s \rightarrow -\infty$.

The logic of scenarios (1) and (2) follow from Proposition 2 in Callander (2011) and the optimal one period behavior in Proposition 1 in this paper. In both scenarios, the second period utility for the Voter is exactly equal to that in the first period. Thus, the Voter's two-period utility approaches his one-period utility as $s \rightarrow -\infty$. The proposition follows by

continuity.

References

- Anesi, V. and T. R. Bowen (2021). Policy experimentation in committees: A case against veto rights under redistributive constraints. *American Economic Journal: Microeconomics* 13(3), 124–62.
- Banks, J. S. and J. Duggan (2000, March). A bargaining model of collective choice. *American Political Science Review* 94(1), 73–88.
- Baron, D. P. (1996). A dynamic theory of collective goods programs. *American Political Science Review* 90(2), 316–330.
- Buisseret, P. and D. Bernhardt (2017). Dynamics of policymaking: Stepping back to leap forward, stepping forward to keep back. *American Journal of Political Science* 61(4), 820–835.
- Callander, S. (2008). A theory of policy expertise. *Quarterly Journal of Political Science* 3(2), 123–140.
- Callander, S. (2011, November). Searching for good policies. *American Political Science Review* 105, 643–662.
- Cameron, C. M. (2000). *Veto Bargaining: Presidents and the Politics of Negative Power*. Cambridge University Press.
- Clinton, J. D. (2012). Congress, lawmaking, and the fair labor standards act, 1971–2000. *American Journal of Political Science* 56(2), 355–372.
- Gieczewski, G. and S. Kosterina (2020). Endogenous experimentation in organizations. *Princeton University Typescript*.
- Gilligan, T. and K. Krehbiel (1987). Collective decision making and standing committees: An informational rationale for restrictive amendment procedures. *Journal of Law, Economics, and Organization* 3(2), 287–335.
- Gittins, J. (1979). Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society B* 41(2), 148–177.

- Groseclose, T. and N. McCarty (2001). The politics of blame: Bargaining before an audience. *American Journal of Political Science*, 100–119.
- Howell, W. G. and T. M. Moe (2016). *Relic: How Our Constitution Undermines Effective Government—and Why We Need a More Powerful Presidency*. Basic Books.
- Howell, W. G. and T. M. Moe (2020). *Presidents, Populism, and the Crisis of Democracy*. University of Chicago Press.
- Kalandrakis, A. (2004). A three-player dynamic majoritarian bargaining game. *Journal of Economic Theory* 116(2), 294–14.
- Krehbiel, K. (1998). *Pivotal Politics: A Theory of US Lawmaking*. University of Chicago Press.
- McCarty, N. (1997). Presidential reputation and the veto. *Economics & Politics* 9(1), 1–26.
- Merton, R. K. (1936). The unanticipated consequences of purposive social action. *American sociological review* 1(6), 894–904.
- Messner, M. and M. K. Polborn (2012). The option to wait in collective decisions and optimal majority rules. *Journal of Public Economics* 96(5-6), 524–540.
- Neumark, D. and P. Shirley (2021). Myth or measurement: What does the new minimum wage research say about minimum wages and job loss in the united states? *Industrial Relations: A Journal of Economy and Society*.
- Penn, E. M. (2009). A model of farsighted voting. *American Journal of Political Science* 53(1), 36–54.
- Primo, D. M. (2002). Rethinking political bargaining: Policymaking with a single proposer. *Journal of Law, Economics, and Organization* 18(2), 411–427.
- Romer, T. and H. Rosenthal (1978). Political resource allocation, controlled agendas, and the status quo. *Public choice*, 27–43.
- Rosenthal, H. and W. R. Zame (2022). Sequential referenda with sophisticated voters. *Journal of Public Economics* 212, 104702.
- Strulovici, B. (2010). Learning while voting: Determinants of collective experimentation. *Econometrica* 78(3), 933–971.

Vaghul, K. and B. Zipperer (2021). Historical state and sub-state minimum wages. Technical Report Version 1.3.0, Murray Research Archive. url <https://github.com/benzipperer/historicalminwage/releases/tag/v1.3.0>.

Woon, J. and I. P. Cook (2015). Competing gridlock models and status quo policies. *Political Analysis* 23(3), 385–399.