

The Dynamic (In)efficiency of Monetary Policy by Committee

Alessandro Riboni and Francisco J. Ruge-Murcia*

January 2006

Abstract

This paper develops a model where the value of the monetary policy instrument is selected by a heterogeneous committee engaged in a dynamic voting game. Committee members differ in their institutional power and, in certain states of nature, they also differ in their preferred instrument value. Preference heterogeneity and concern for the future interact to generate decisions that are dynamically inefficient and inertial around the previously-agreed instrument value. This model endogenously generates autocorrelation in the policy variable and provides an explanation for the empirical observation that the nominal interest rate under the central bank's control is infrequently adjusted.

JEL Classification: E58, D02

Key Words: Committees, status-quo bias, interest-rate smoothing, dynamic voting

*Both authors: Département de sciences économiques, Université de Montréal. We thank participants in the conference on New Developments in the Analysis of Monetary Policy and Institutions at Tel Aviv University (December, 2005) for helpful comments. This research received the financial support of the Social Sciences and Humanities Research Council of Canada. *Correspondence:* Francisco J. Ruge-Murcia, Département de sciences économiques, Université de Montréal, C.P. 6128, succursale Centre-ville, Montréal (Québec) H3C 3J7, Canada. E-mail: francisco.ruge-murcia@umontreal.ca

1 Introduction

This paper studies the dynamic implications of monetary policy making by committee. The subject matter is important because, in many countries, monetary policy decisions are made by committees, rather than by one individual alone. For example, Fry *et al.* (2000) report that in a sample of 88 central banks, 79 use some form of committee structure to formulate monetary policy.

In particular, this paper focuses on a two-person committee where heterogenous agents must select the value of the policy instrument (say, the nominal interest rate) but face exogenous uncertainty regarding their preferred policies in the future. The committee members differ in two ways. First, agents have different state-dependent preferences over policy. There are states of nature where agents do not agree in their preferred instrument value, and states where they agree. Second, agents differ in their institutional role. More concretely, one agent, the chairman or agenda setter, makes a take-it-or-leave-it proposal to the other agent in every period. This assumption captures the idea that chairmen usually have more power and influence than their peers as a result of additional legal responsibilities, statutory prerogatives, or prestige. The identity of the chairman and the composition of the committee are assumed to be fixed over time. An important and plausible feature of the voting game is that the instrument value decided in the previous meeting is the default option in case the proposal is rejected in the current meeting. Hence, the current status quo is a state variable.¹

In this setup, the first-best policy (that is, the state-contingent program that a benevolent social planner would choose) prescribes the policy preferred by both agents in states of agreement and optimal risk-sharing in states of disagreement. However, since the implementation of this optimal plan requires commitment, it is not surprising that the politico-economic equilibrium cannot implement the first-best policy in the absence of a commitment technology. Instead, the politico-economic equilibrium features inefficient policy choices in all states of nature. First, in states of agreement, committee members do not select their common preferred policy. The reason is that forward-looking policy makers realize that current decisions affect future voting outcomes by changing the default option in the next meeting. Hence, in choosing the current policy, committee members trade-off the benefit of selecting their preferred policy for this period and the cost of affecting their bargaining power in future states of disagreement. Simulations show that this form of “political failure” (to borrow

¹The theoretical literature on bargaining with evolving defaults is scant. Among the few contribution, see Baron (1996), Baron and Herron (2003), and Bernheim *et al.* (2005). To the best of our knowledge, this paper is the first to incorporate preference uncertainty in a model with dynamic reversion of the status quo.

the term proposed by Besley and Coate, 1998) often implies muted responses to changing economic conditions and, consequently, provides a rationale for policy conservatism.²

Second, in states of disagreement, committee decisions are inertial. That is, the chairman's optimal proposal is often the status quo, even when the state of nature has changed. This result is primarily due to the heterogeneity in policy preferences and to the role of the status quo as the default option in the voting game. Since the default policy may not be undesirable for a committee member, in many instances policy changes are not passed (or proposed). A similar status-quo bias in policy making is derived by Romer and Rosenthal (1978) in a static model. However, compared to their agenda-setting game, the status-quo bias in this model gives rise to inefficiency. In particular, there is imperfect risk-sharing in that there are instances where the current policy is (close to) optimal for one policy maker but very costly for the other (or *vice versa*). Also, the status-quo bias in this model is not as severe as in the game in Romer and Rosenthal. In some circumstances, the chairman is able to change policy even if the default option coincides with the preferred policy of the other member. This result arises because, in a dynamic setup, committee members smooth their bargaining power across states and are willing to lower their current utility to increase their bargaining power in future meetings.

The inertia in committee decision making predicted by this model is a plausible explanation for the empirical observation that the interest rate under a central bank's control is infrequently adjusted despite the arrival of new information. Figure 1 plots the histogram of the changes in the target value of the key interest rate in four central banks: the U.S. Federal Reserve, the European Central Bank, the Bank of England, and the Bank of Canada.³ Note that, by far, the most frequent policy decision is to leave the interest-rate target unchanged. In addition, since the status quo is a state variable, current and lagged instrument values are linked through the solution of the chairman's dynamic optimization problem. Hence, this model endogenously generates autocorrelation in key interest rates. In contrast, the standard model with a single central banker, which underlies the derivation of the Taylor rule, predicts that the interest rate is always adjusted whenever economic conditions change and

²Other explanations for partial adjustment to shocks include policy maker uncertainty (Orphanides, 2003), improved control over long-term interest rates (Goodfriend, 1991, and Woodford, 2003), and reduction of financial stress (Cukierman, 1991).

³The interest rates are the Federal Funds Rate, the Rate for Main Refinancing Operations, the Repo Rate, and the Overnight Rate, respectively. The samples used to construct these histograms start in August 1987, January 1999, June 1997, and June 1997, respectively, and end in March 2005 in all cases. For the Federal Reserve, the sample starts with the first meeting under Alan Greenspan chairmanship, and the data sources are Chappell *et al.* (2005) and the Federal Reserve Bank of New York. For the other central banks, the sample starts (roughly) at the time when committee decision making was instituted, and the data were collected by the authors using official press releases.

does not predict interest rate autocorrelation.⁴ Since interest rates are serially correlated in the data, lagged interest rates are usually appended to the Taylor rule in empirical work.

As a result of policy inertia, the policy variable in this model changes less often than the state of nature and, consequently, the path of the former is smoother than that of the latter. A similar result whereby committee decision making induces policy smoothing has been derived by Waller (2000) in a model with partisan central bank appointments and exogenous electoral outcomes *à la* Alesina (1987). In our model, policy smoothing is not sustained by the strategic appointment of moderate committee members (as in Waller’s model) or by trigger punishments (as in Alesina’s model), but is instead the result of the voting game played by the committee. Moreover, in the above literature, policy smoothing is regarded as welfare increasing because it reduces the uncertainty associated with elections. Thus, a constant policy rule, irrespective of the identity of the winning party, is beneficial to both parties. In our model, preferred policies are not constant but instead vary over time as the state of nature changes. As a result, a constant policy is not optimal and policy inertia moves the economy away from the efficient frontier.

The paper is organized as follows. Section 2 describes the committee and solves a simple two-state model that illustrates the main implications of the voting game. Section 3 solves and simulates a more general multi-state model. Section 4 compares the voting model with an endogenous and a fixed default. Section 5 concludes.

2 Two-State Model

This section describes the committee and examines a version of the dynamic voting game with two states of nature. The two-state model is solved for three horizons, namely $T = 1, 2$ and ∞ . The finite horizon cases ($T = 1, 2$) are solved analytically by backward induction and the infinite horizon case ($T = \infty$) is solved numerically. Studying the two-state model first helps develop the reader’s intuition by illustrating some of our results in the simplest possible setup.

The committee is composed of two agents with heterogenous preferences: C and P , where C is the fixed chairman.⁵ In every period, the committee is concerned with selecting the policy variable x that takes values in the interval $[a, c]$, with $a < c$. To make this more

⁴An exception is Woodford (2003) where a motive for interest-rate smoothing is explicitly introduced into the central bank’s objective function.

⁵The assumption of a fixed agenda setter is made for the sake of realism. For example, in the case of the United States, the chairman of the Federal Open Market Committee is (by tradition) the chairman of the Board of Governors, who in turn is appointed by the President for a renewable four-year term. For models of legislative bargaining where the the agenda setter is randomly selected, see Baron (1996) and Baron and Ferjohn (1989).

concrete, think of the policy variable as the target value of a key nominal interest rate. In each period, the payoff of policy maker j , for $j = C, P$, is

$$U_j(x, \varepsilon) = -(x - r_j(\varepsilon))^2,$$

where $r_j(\varepsilon)$ is j 's state-dependent preferred policy and ε is an exogenous shock. For analytical convenience, it is assumed that the probability distribution of ε is discrete. In this section, it is also assumed that ε can take only two values, ε_1 and ε_2 . The shock follows a Markov chain and its transition matrix has elements $p_{ki} = \text{prob}(\varepsilon_k | \varepsilon_i) \in (0, 1)$ with $i, k = 1, 2$ and $\sum_{k=1}^2 p_{ki} = 1$. Two states of nature are defined by the possible realizations of ε . When $\varepsilon = \varepsilon_1$, agents C and P disagree in their preferred instrument values, with c and a their respective preferred points.⁶ When $\varepsilon = \varepsilon_2$, C and P agree and $b \in (a, c)$ is their preferred point. Without loss of generality, it is assumed that the bliss points are evenly spaced, meaning that $b - a = c - b$.

Each committee member ranks policy sequences according to the expected utility they deliver. The intertemporal utility of member j is

$$E \left(\sum_{t=1}^T \delta^{t-1} U_j(x_t, \varepsilon_t) \right),$$

where $\delta \in (0, 1)$ is the discount factor, which is the same for both players. Note that preferences depend on the policy instrument rather than on policy outcomes (say, inflation and unemployment). This approach has two advantages. First, it makes the voting game more tractable because otherwise the private sector's expectations would be a state variable that has to be validated in a rational expectations equilibrium.⁷ Second, it means that the particular economic model that the policy maker believes to be true need not be specified. This is important because anecdotal evidence suggests that policy makers may have different views about how the economy works depending on their background and intellectual environment. For example, Hetzel (1998) argues that Arthur Burns, the chairman of the U.S. Federal Reserve from February 1970 to January 1978, attached significant importance to nonmonetary factors (*e.g.*, business optimism and wage demands by unions) in the determination of inflation, and did not consider monetary policy to be the driving force behind

⁶The converse assumption – that C and P 's preferred points are a and c , respectively – leads to decision rules that are mirror images of the ones derived here. Hence, the main theoretical implications of the model are robust to using either version of this assumption.

⁷This is a non-trivial fixed point to solve for. The strategy of the private sector depends on the expected voting outcome, but the outcome of the voting game depends on the expectations of the private sector in two ways: 1) directly, because expected inflation affects policy makers' utilities; and 2) indirectly by changing the default payoff, since the real interest rate in case of disagreement is the difference between the nominal status quo policy and expected inflation. For now, we leave this extension to future research.

the rise in U.S. inflation in the 1970s. The actions of Paul Volcker as chairman from August 1979 to August 1987 suggests that he did not completely share Burns' views on the causes of inflation. Furthermore, committee members serving under the same chairman may disagree on the "correct" economic model. For example, Chappell *et al.* (2005, ch. 6.3) document the division within the Federal Open Market Committee (FOMC) between Keynesian and Monetarist members during Burns' chairmanship.

Before discussing how the committee makes decisions, we derive the benchmark first-best. Let x_i^* denote the optimal policy when the state of nature i has occurred, with $i = 1, 2$. The first-best (x_1^*, x_2^*) is given by (x, b) , where x is any policy in the interval $[a, c]$. The fact that $x_2^* = b$ is obvious. To see that x_1^* can be any instrument value in the interval $[a, c]$, recall that members' preferences are opposite when $\varepsilon = \varepsilon_1$ takes place and, consequently, it is not possible to Pareto-improve upon any $x \in [a, c]$.

The committee decides policies sequentially with the following timing. First, the current realization of the shock ε is observed. Then, the chairman makes a take-it-or-leave-it proposal $x \in [a, c]$. If the proposal is rejected by P , then the status quo persists until the next period. If the proposal is accepted, then x is implemented and becomes the new status quo for the voting game in the next period. The assumption that the chairman makes take-it-or-leave-it proposals to the committee is not meant to be a literal description of how monetary committees actually work. Instead, it is a modeling device that captures the idea that chairmen usually have more power and influence than their peers.⁸

Mathematically, the problem of the chairman can be formulated recursively with state given by the initial status quo and the current shock. The Markov strategies of the two agents are defined as follows. The proposal strategy of the chairman is

$$G_{C,t} : [a, c] \times \{\varepsilon_1, \varepsilon_2\} \rightarrow [a, c].$$

The voting rule followed by P depends on both the state and the proposal made by C ,

$$G_{P,t} : [a, c] \times \{\varepsilon_1, \varepsilon_2\} \times [a, c] \rightarrow \{yes, no\}.$$

In the infinite-horizon game, player's strategies will be stationary. The voting rule is assumed to be sequentially rational. That is, P votes in favor of the proposal whenever the current utility from the proposal plus the continuation value of moving to the next period with a

⁸On the account of his experience as Board governor from 1996 to 2002, Laurence Meyer (Meyer, 2004) remarks on "the chairman's disproportionate influence on FOMC decisions" and on "his efforts to build consensus around his policy recommendations" (p. 50). However, Mayer also notes that the chairman "does not necessarily always get his way" (p. 52). Sherman Maisel, who was member of the Board during Burns' chairmanship also points out that "while the influence of the Chairman is indeed great, he does not make policy alone" (Maisel, 1973, p. 124).

new status quo is higher than or equal to keeping the status quo and moving to the next period with the current status quo. Define the acceptance set A_t as the set of policies that are acceptable by P at time t , for a given default policy and a given realization of the state of nature. More formally,

$$A_t(q, \varepsilon_i) = \left\{ x \in [a, c] : U_P(x, \varepsilon_i) + \delta \sum_{k=1}^2 p_{ki} V_{P,t+1}(x, \varepsilon_k) \geq U_P(q, \varepsilon_i) + \delta \sum_{k=1}^2 p_{ki} V_{P,t+1}(q, \varepsilon_k) \right\},$$

where the sum term is the conditional expectation of the value function $V_{P,t+1}(\cdot, \varepsilon)$ as of time t . Note that unanimity is required for a policy change only because we are considering a two-person committee. Appendix A shows that our setup is equivalent to a committee with $n + 1$ representatives where P occupies the role of the median and a simple majority is required to pass a proposal.

Let $q \in [a, c]$ denote the initial status quo. For all t the proposal strategy $G_{C,t}(q, \varepsilon_i)$ solves the dynamic programming problem

$$V_{C,t}(q, \varepsilon_i) = \max_{x \in A_t(q, \varepsilon_i)} \left(U_C(x, \varepsilon_i) + \delta \sum_{k=1}^2 p_{ki} V_{C,t+1}(x, \varepsilon_k) \right).$$

In words, C proposes the policy x that maximizes his utility from among those that are acceptable to P . In the noncooperative bargaining environment studied here, the chairman's proposals are never rejected in equilibrium.⁹ The latter implication is in line with historical records from the FOMC which show that a chairman's recommendation has never been voted down by the committee (see, Chappell *et al.*, 2005).

The Markov perfect equilibrium of this game is a set of policy rules $\{G_{C,t}, G_{P,t}\}_{t=1}^T$, such that: 1) for all t the voting rule $G_{P,t}$ is sequentially rational given $\{G_{P,s}\}_{s=t+1}^T$ and $\{G_{C,s}\}_{s=t}^T$; and 2) for all t the proposal rule $G_{C,t}$ solves the problem of the agenda setter at time t , given $\{G_{C,s}\}_{s=t+1}^T$ and $\{G_{P,s}\}_{s=t}^T$.

2.1 Finite Horizon with $T=1$

Consider the voting game described above with finite horizon $T = 1$. Absent any dynamics, the solution is similar to that of the agenda-setting game studied by Romer and Rosenthal (1978). The chairman's proposal strategy is depicted in the first column of Figure 2 as a function of the status quo q for each possible realization of ε . Proposals on the 45 degree line are the status quo.

⁹Note that proposing a policy outside the acceptance set is equivalent to proposing the status quo, which is always accepted.

First, suppose that ε_1 occurs. In this case, the chairman proposes the status quo for any $q \in [a, c]$. The reason is that P would not accept any proposal $x \in (q, c]$ that gives C higher utility than q , and C would not propose any $x \in [a, q)$ that gives him lower utility than q . Since the proposal strategy is independent of the values of δ, p_{11} and p_{22} , it follows that policy inertia arises in this case only as a result of the heterogeneity among committee members. Now, suppose that ε_2 occurs and both members agree that b is the optimal value of the policy instrument. In this case, the chairman proposes b starting from any status quo. Notice that the outcome of this (static) game coincides with the first-best.

2.2 Finite Horizon with $T=2$

Suppose now that the horizon is $T = 2$. The model is solved backwards for $t = T, T - 1$. The proposal strategies at time $t = T$ are the ones derived in Section 2.1. The proposal strategies at time $T - 1$ are derived in Proposition 1 below. In order to develop the reader's intuition, these strategies are depicted in the second column of Figure 2 in the special case where $\delta = 0.5$, $p_{11} = 0.8$ and $p_{22} = 0.5$. These probabilities correspond approximately to those computed using the voting records of the Monetary Policy Committee (MPC) of the Bank of England from June 1997 to January 2005.¹⁰

Proposition 1. *Let $T = 2$. For all $q \in [a, c]$ the proposal rules at time $T - 1$ when ε_1 and ε_2 occur are, respectively, $G_{C,T-1}(q, \varepsilon_1) = q$ and*

$$G_{C,T-1}(q, \varepsilon_2) = \begin{cases} y, & \text{for } q \in [a, 2z - y], \text{ where } y = (b + c\delta p_{12})/(1 + \delta p_{12}), \\ 2z - q, & \text{for } q \in (2z - y, z), \text{ where } z = (b + a\delta p_{12})/(1 + \delta p_{12}), \\ q, & \text{for } q \in [z, y], \\ y, & \text{for } q \in (y, c]. \end{cases}$$

Proof: We start by showing that $G_{C,T-1}(q, \varepsilon_1) = q$ is the optimal proposal rule. Suppose that the current shock is ε_1 . The chairman's proposal strategy at time $t = T - 1$ is found by exploiting the fact that the successful proposal in T will be given by the proposal rules in Section 2.1. The chairman chooses the proposal x that maximizes his two-period payoff

¹⁰The voting records contain information on: the date of the meeting; the policy decision; the names of members in favor of the decision; and the names and preferred policy options of dissenting members. The probabilities are computed as follows. A meeting where the policy decision is adopted unanimously is treated as one where all committee members agree in their preferred instrument value, meaning that in terms of our model $\varepsilon = \varepsilon_2$. A meeting with at least one dissenting individual is treated as one where committee members disagree in their preferred instrument value, meaning that $\varepsilon = \varepsilon_1$. Then, p_{11} (p_{22}) is computed as the number of observations where members disagree (agree) in two consecutive meetings divided by the number of observations where members disagree (agree) in the first of these two meetings. Since the mapping from the voting records to the model is clearly imperfect, the policy rules in Figure 2 are best interpreted as illustrative only.

within the acceptance set, $A_{T-1}(q, \varepsilon_1)$. That is, he solves the following problem:

$$\max_{x \in A_{T-1}(q, \varepsilon_1)} -(1 + \delta p_{11})(x - c)^2,$$

where the acceptance set is defined as

$$A_{T-1}(q, \varepsilon_1) = \{x \in [a, c] : -(1 + \delta p_{11})(x - a)^2 \geq -(1 + \delta p_{11})(q - a)^2\}.$$

It is easy to see that the acceptance set is $[a, q]$ for any $q \in [a, c]$. Since C 's two-period payoff is increasing in the current proposal, the chairman always proposes $x = q$.

Now we prove that the posited $G_{C, T-1}(q, \varepsilon_2)$ is optimal. When ε_2 occurs at time $T - 1$, the chairman's problem becomes:

$$\max_{x \in A_{T-1}(q, \varepsilon_2)} -(x - b)^2 - \delta p_{12}(x - c)^2,$$

where

$$A_{T-1}(q, \varepsilon_2) = \{x \in [a, c] : -(x - b)^2 - \delta p_{12}(x - a)^2 \geq -(q - b)^2 - \delta p_{12}(q - a)^2\}.$$

In finding $G_{C, T-1}(q, \varepsilon_2)$, it is useful to first derive P 's voting rules. P 's two-period utility is concave in x , with a maximum at

$$z = \frac{b + a\delta p_{12}}{1 + \delta p_{12}}.$$

Note that $a < z < b$. Because the payoff is symmetric around z , the acceptance set is easy to derive. For any $q \in [a, z]$, $A_{T-1}(q, \varepsilon_2) = [q, 2z - q]$, and for any $q \in [z, c]$, $A_{T-1}(q, \varepsilon_2) = [2z - q, q]$. Now consider C 's proposal strategy. C 's objective function is concave and has a global maximum at

$$y = \frac{b + c\delta p_{12}}{1 + \delta p_{12}}.$$

Note that $b < y < c$. When $q \in [y, c]$, C is not constrained and will propose y . When $q \in (2z - y, y)$, C is constrained and proposes his preferred policy in the acceptance set. We distinguish two cases: when $q \in [z, y)$, the proposal is $x = q$, and when $q \in (2z - y, z)$, the proposal is $x = 2z - q$. Finally, when $q \in [a, 2z - y]$, the acceptance set includes C 's bliss point y and, consequently, C proposes $x = y$. ■

Note that the decision rules in period $T - 1$ converge to those in period T as $\delta \rightarrow 0$ (committee members attach no weight to future payoffs) or $p_{11}, p_{22} \rightarrow 1$ (the states of nature are absorbing): in either case $y, z \rightarrow b$.

We now comment on the policy rules just derived. When ε_1 occurs and committee members disagree on the optimal instrument value, the proposal at time $T - 1$ is q , irrespective

of the current status quo. The status-quo bias originates from the opposite preferences of the two players. In this case, there is no Pareto-improving policy change and the political equilibrium is efficient according to the standard economic definition. To see this, pick any $q \in [a, c]$ and note that any policy choice to the right (left) of q would reduce P 's (C 's) utility. Thus, for the two-state model, the committee implements the first-best in the state of disagreement.¹¹

When ε_2 occurs and committee members agree that b is the optimal instrument value today, Proposition 1 shows that the proposal at time $T - 1$ is generically different from the first-best policy b . For example, when $q \in [y, c]$, C adjusts the current policy only to y , which is larger than b . This result is due to the non-zero probability of disagreement in the next meeting: the chairman trades off the benefit of moving towards the ideal point b and the cost of weakening his bargaining power should $\varepsilon = \varepsilon_1$ in the next period.

To verify the existence of a political failure in equilibrium, consider, for instance, the case where $q = y$. Rather than staying in y , as established in Proposition 1, a Pareto-improving choice would be moving to b today and going back to y in the next period should ε_1 occur. However, this policy requires commitment. Absent commitment, after the default policy has changed to b , it is not sequentially rational for P to allow C to return to y . Consequently, a policy change to b will not be implemented by the committee.¹² Similar policies that Pareto-improve upon those in Proposition 1 can be constructed for all status quo in $[a, c]$ except for $\{b, 2z - b\}$. (In these two cases, committee decision making implements the first-best because the proposal coincides with b .) Hence, this simple two-state, two-period model illustrates the fact that, in some circumstances, committee decision making with an endogenous status quo is inefficient; responses to shocks are more muted compared to situations where there is a single central banker. Since y is increasing in p_{12} and δ , the chairman becomes more cautious as the conditional probability of future disagreement increases and as the future is discounted less heavily by committee members.

While it is difficult to obtain a complete characterization of the proposal rule for an arbitrary period $T - s$, where s denotes the number of remaining periods until T , the result of partial adjustment carries over as T increases. Suppose the status quo at time $T - s$ is equal to c . We can show that the committee does not move when ε_1 occurs and moves to

¹¹Section 3 below shows that this result is not robust to increasing the number of shock realizations and, consequently, in the multi-state version of the model, the politico-economic equilibrium is inefficient in all states.

¹²This model abstracts from sunset proposals, that is proposals over more than one period. However, if sunset proposals are allowed while keeping the assumption that the default in the next meeting is the status quo, then our results would be unchanged.

v_{T-s} when ε_2 occurs, where v_{T-s} is defined as

$$v_{T-s} = \frac{b + c\delta p_{12} \sum_{j=1}^s (\delta p_{11})^{j-1}}{1 + \delta p_{12} \sum_{j=i}^s (\delta p_{11})^{j-1}}, \quad 1 \leq s \leq T - 1.$$

Note that $v_{T-s} = y$ in the special case where $T = 2$ and $s = 1$. If repeated realizations of $\varepsilon = \varepsilon_2$ take place, the committee moves gradually towards $y > b$. To see this, note that the sequence $\{v_{T-s}\}_{s=1}^{T-1}$ is increasing in s and converges to y as the economy approaches the previous-to-last period. Intuitively, at time $T - s - 1$, the chairman is more cautious in moving towards y than at $T - s$ because he is more likely to be constrained as a result of the current choice when there are more periods left before the end of the game. Note that today's decision has an effect on future outcomes only when ε_1 occurs in the next period, two periods in a row, three periods in a row, etc.¹³ However, when there are more periods left before the end of the game, the sum of the probabilities associated with these events is quantitatively larger.

2.3 Infinite Horizon

Consider now the voting game in the case where the horizon is infinite. Because finding the analytical solution to the infinite-horizon game is not trivial, we employ instead a numerical algorithm to find the stationary decision rules. The procedure builds on the projection method employed by Judd (1998) to study the Bellman equation of the stochastic growth model, and works by backward induction exploiting the observation that the chairman's problem is a constrained maximization which can be solved numerically using standard hill-climbing methods. See Appendix B for a detailed description of the algorithm.

The chairman's stationary decision rules are plotted in the third column of Figure 2. When $\varepsilon = \varepsilon_1$ and both members disagree, the chairman simply proposes the status quo. When $\varepsilon = \varepsilon_2$, the proposal strategy is qualitatively similar to that derived analytically in Proposition 1 for the horizon $T = 2$, but the difference between the proposed policy and the current bliss point is larger. This result follows from the observation that v_{T-s} decreases as $s \rightarrow \infty$. In particular,

$$\lim_{s \rightarrow \infty} v_{T-s} = v = \frac{b + c\delta p_{12}/(1 - \delta p_{11})}{1 + \delta p_{12}/(1 - \delta p_{11})} > \frac{b + c\delta p_{12}}{1 + \delta p_{12}} = y > b.$$

¹³For example, if ε_1 occurs in the next period and ε_2 in two periods from now, the payoff two periods from now is not affected by the current decision because two periods from now the chairman will not be constrained and can move to his bliss point.

Finally, since $v - w > y - z$ (see Figure 2), the set of status quo for which the chairman does not propose a policy change is larger in the infinite horizon case.

3 Multi-State Model

This section solves the dynamic voting game in the more general case where the number of possible shock realizations is larger than two. This extension is important for two reasons. First, it shows that the efficient outcome in the state of disagreement reported in Section 2 is not robust to increasing the number of shock realizations and, consequently, committee policy choices may be inefficient in all states. Second, the two-state model features a strong form of policy inertia in the form of the absorbing region $[w, v]$ and, consequently, it does not permit the derivation of time series implications.¹⁴ In what follows, the chairman's proposal strategies are computed, and then policy decisions by the committee are simulated for a sample of sequential meetings.

3.1 Proposal Strategies

Assume that the shock ε can take I discrete values, ε_i for $i = 1, 2, \dots, I$. Define $S = \{\varepsilon_1, \varepsilon_2, \dots, \varepsilon_I\}$. As before, the shock follows a Markov chain and its $I \times I$ transition matrix has elements $p_{ki} = \text{prob}(\varepsilon_k | \varepsilon_i) \in (0, 1)$ that satisfy $\sum_{k=1}^I p_{ki} = 1$. The shock ε shifts the agents' preferred policies over a policy set denoted by X . The timing and other features of the model are as described in Section 2. For this more general specification, the Markov strategies of the two agents are defined by

$$\begin{aligned} G_{C,t} &: X \times S \rightarrow X, \\ G_{P,t} &: X \times S \times X \rightarrow \{yes, no\}. \end{aligned}$$

The chairman's proposal strategy $G_{C,t}(q, \varepsilon_i)$ solves the dynamic programming problem

$$V_{C,t}(q, \varepsilon_i) = \max_{x \in A_t(q, \varepsilon_i)} U_C(x, \varepsilon_i) + \delta E_t V_{C,t+1}(x, \varepsilon),$$

where E_t denotes the conditional expectation at time t and the acceptance set is defined as

$$A_t(q, \varepsilon_i) = \{x \in X : U_P(x, \varepsilon_i) + \delta E_t V_{P,t+1}(x, \varepsilon) \geq U_P(q, \varepsilon_i) + \delta E_t V_{P,t+1}(q, \varepsilon)\}.$$

For concreteness, we focus on the case where $I = 6$ and maintain the convention that committee members agree in the even states and disagree in the odd states of nature. The

¹⁴To see this, note that as soon as ε_2 occurs, the successful proposal will be $x \in [w, v]$ with $x = q \in [w, v]$ thereafter.

bliss points of P (C) in states 1 through 6 are, respectively, $a(c)$, $b(b)$, $b(d)$, $c(c)$, $c(e)$, and $d(d)$, where $a < b < c < d < e$ and are equally spaced. The policy set is the interval $[a, e]$. Stationary decision rules are solved for using the algorithm described in Section 2.3.

In what follows, we characterize the first-best policy for the multi-state version of the model. As before, x_i^* denotes the first-best policy when shock i occurs, with $i = 1, \dots, 6$. Clearly, $x_2^* = b$, $x_4^* = c$, and $x_6^* = d$. Regarding the optimal policies in the odd states, the following risk-sharing conditions must hold:

$$\frac{U'_P(x_i^*, \varepsilon_i)}{U'_C(x_i^*, \varepsilon_i)} = \kappa, \quad \text{for } i = 1, 3, 5,$$

where $\kappa < 0$ is a constant. That is, the ratio of marginal utilities is equalized across all states of disagreement. For the functional form of the payoff function used here,

$$\frac{x_1^* - a}{x_1^* - c} = \frac{x_3^* - b}{x_3^* - d} = \frac{x_5^* - c}{x_5^* - e} = \kappa.$$

From this condition it follows that $x_1^* - a = x_3^* - b = x_5^* - c$, where $a \leq x_1^* \leq c$, $b \leq x_3^* \leq d$, and $c \leq x_5^* \leq e$. To see, for example, that $x_1^* - a = x_3^* - b$, suppose, on the contrary, that $x_1^* - a > x_3^* - b$ ($x_1^* - a < x_3^* - b$) and note that in this case both policy makers could augment their payoff by lowering (increasing) x_1^* and increasing (decreasing) x_3^* .

Since the chairman's proposal strategies depend on the matrix of transition probabilities, we conducted extensive experiments with various parameter configurations and report below results for $\delta = 0.5$ and the transition matrices¹⁵

$$A = \begin{bmatrix} 3/5 & 1/5 & 0 & 0 & 0 & 0 \\ 1/5 & 3/5 & 1/5 & 0 & 0 & 0 \\ 1/5 & 1/5 & 3/5 & 1/5 & 0 & 0 \\ 0 & 0 & 1/5 & 3/5 & 1/5 & 1/5 \\ 0 & 0 & 0 & 1/5 & 3/5 & 1/5 \\ 0 & 0 & 0 & 0 & 1/5 & 3/5 \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} 1/6 & 1/6 & 1/6 & 1/6 & 1/6 & 1/6 \\ 1/6 & 1/6 & 1/6 & 1/6 & 1/6 & 1/6 \\ 1/6 & 1/6 & 1/6 & 1/6 & 1/6 & 1/6 \\ 1/6 & 1/6 & 1/6 & 1/6 & 1/6 & 1/6 \\ 1/6 & 1/6 & 1/6 & 1/6 & 1/6 & 1/6 \\ 1/6 & 1/6 & 1/6 & 1/6 & 1/6 & 1/6 \end{bmatrix}.$$

Matrix A was deliberately designed to represent the idea that preferred policies evolve slowly over time as new information about business cycle and inflation variables becomes available. Matrix B is used to show that: 1) committee decision making can generate endogenous autocorrelation in the policy variable even when the states of nature are not autocorrelated; and 2) overshooting may be an outcome of the voting model. Decision rules are respectively plotted in the first and second column of Figure 3. Proposals on the 45 degree line are the status quo (that is, $x = q$).

¹⁵The relatively low value of δ is used to show that dynamic inefficiency arises in the multi-state version of the model even when the future is heavily discounted. Alesina (1987) argues that policy makers' effective discount rates may be low because reappointment probabilities are less than one. Results from unreported experiments are available from the corresponding author upon request.

The following implications for committee decision making can be drawn from Figure 3. First, consider the proposal rules in states of agreement. As before, the chairman proposes instrument values different from b, c and d in states 2, 4 and 6, respectively, even though both members agree that these are their current preferred policy options. The intuition for this result is the same as in the two-state model, namely that in a dynamic setup, committee members face a trade-off between the current benefit of choosing their preferred policy and the possible cost of reducing their bargaining power in future meetings. In most cases, the committee partially adjust to shocks that align preferences and, consequently, policy changes are typically smaller than the optimal ones. While policy conservatism is by far the most common outcome, overshooting may arise when drastic changes in the preferred policies are allowed. By overshooting, we refer to the situation where the committee changes the instrument value by more than a single central banker would. Then, policy changes are larger than the optimal ones. An example of overshooting under the transition matrix B is the following. Starting in state $\varepsilon = \varepsilon_2$ and with a status quo larger than b , note that the chairman proposes a policy less than b , while the single central banker would have adopted b .¹⁶ Note that, like partial adjustment, overshooting is also inefficient because both committee members would increase their current payoff by choosing the instrument value they currently prefer.

Now consider the proposal rules in states of disagreement. In these cases, there is local policy inertia around previously agreed on decisions. To see this, consider the following example. Starting from state $\varepsilon = \varepsilon_2$ and instrument value b , suppose there is a “small” change in the state of nature, meaning to either of the adjacent states $\varepsilon = \varepsilon_1$ or ε_3 . In these states, members disagree on their preferred instrument value but the chairman’s decision rule still implies $x = b$. Now, suppose there is a “large” change in the state of nature, meaning to $\varepsilon = \varepsilon_4, \varepsilon_5$ or ε_6 . Note that in these cases the proposal will be different from the status quo regardless of whether members agree in their desired instrument value or not. An implication of local inertia is that the relation between changes in the state of nature and in policy is nonlinear. In particular, small changes in the state of nature are less likely to produce policy changes compared with larger ones. Empirically, this would mean, for example, that small variations in the rates of inflation and unemployment are less likely to result in a change in the key nominal interest rate, compared with large movements

¹⁶The reason why we observe overshooting with matrix B , but not with matrix A , is the following. The rationale for overshooting and proposing a policy less than b is to have more leverage should $\varepsilon = \varepsilon_3$ occur and get closer to the ideal point e . The cost of overshooting is that the chairman is worse off if shock ε_1 occurs, because the agenda setter is stuck with a policy lower than b , when his ideal instrument value is c . Since $p_{52} = 0$ in matrix A , the expected cost of overshooting is larger, and, consequently, overshooting does not occur in equilibrium.

in these variables.¹⁷ In contrast, the standard model with a single central banker, which underlies the derivation of the linear Taylor rule, predicts a proportional change in the policy instrument for any change in inflation and unemployment regardless of their size.

Note that P allows a policy change in the (odd) states of nature where there is disagreement, even when the current default coincides with his preferred policy. For example, when $q = a$ and $\varepsilon = \varepsilon_1$ occurs, the committee chooses an instrument value closer to c . When the default coincides with his preferred policy, P has significant bargaining power in the current period and, consequently, is willing to accept a policy change to increase his bargaining power in future meetings.¹⁸ This result is not present in the static agenda-setting game of Romer and Rosenthal (1978). It can only be obtained in a dynamic setup where agents have an incentive to smooth their bargaining power across states by choosing the default for the next meeting. This opportunity is valuable because agents are risk-averse. In absence of commitment, agents strategically modify the (endogenous) default in order to better share risk across states. Clearly, this instrument is imperfect: compared to what is prescribed by the first-best, risk-sharing is not optimal (i.e., the politico-economic equilibrium fails to satisfy the efficiency condition in states of disagreement as well). In some states, one of the two policy makers obtains a high payoff while the other suffers a large loss; in some other states, the situation may be reversed. Consequently, there is room for better risk-sharing among committee members.

3.2 Simulations

This section simulates committee decision making using an artificial sample of sequential meetings under the multi-state voting model examined above. This exercise is important because it reveals the proposal strategies that are implemented in practice and permits the derivation of time series implications.

A series of 200 realizations of the shock ε were generated using each transition probability matrix (whether A or B). Then, the outcome of the voting game was found using the chairman's proposal strategies in Figure 3. The simulated series of ε and x are plotted in Figure 4. Notice that there is policy smoothing in the sense that the policy variable changes less often than the state of nature. That is, there are many instances where nature changes but the value of the policy variable remains the same. Earlier research by Alesina

¹⁷Eijffinger, Schaling and Verhagen (1999) construct a model for a single central banker that generates a similar prediction in the form of an inaction range around the previous policy choice, but inertia is the result of an unspecified fixed cost for policy changes.

¹⁸Returning to the previous example, notice that were policy a to remain the default and should $\varepsilon = \varepsilon_5$ occur next period, P would enter the next meeting with very low bargaining power.

(1987) and Waller (2000) also finds that policy may display less variance when decisions are made through committees than when they are made by a single individual. However, in this model, policy smoothing is not sustained by the strategic appointment of moderate committee members (as in Waller’s model) or by trigger punishments (as in Alesina’s model), but by the voting game played by the heterogenous committee. Also, notice that the ergodic process of the policy variable involves a finite number of realizations but they do not correspond to the agreement values (b , c , and d) because of dynamic inefficiency.

From the simulated series, it is possible to construct the frequency histograms for Δx in Figure 5.¹⁹ From this Figure, it is clear that the most common policy decision by the committee is to set $\Delta x = 0$ despite the fact that the state of nature has changed.²⁰ This result is due to the local inertia implied by the optimal decision rules of committee members which was discussed above. Thus, the voting model can provide an explanation for the observation in Figure 1 whereby the interest rate under the central bank’s control is infrequently adjusted, despite the fact that there is new information.

It is important to compare this implication with the one obtained when monetary policy is determined by a single individual, say C . Absent a committee, C ’s decision rule involves changing the policy variable to his preferred value whenever there is a change in the state of nature. The histograms for this case are plotted in the bottom panel of Figure 5 and show that, in contrast with the data, the outcome $\Delta x = 0$ is relatively infrequent.

Figure 6 plots the sample autocorrelation of the policy variable in the model and in the key interest rate from four central banks. First, note that the model endogenously generates positive autocorrelation in the policy variable even when the states of nature are not serially correlated (Matrix B). Second, when the states of nature are persistent (Matrix A), then the predicted autocorrelation may approach that observed in actual data.²¹ Instead, the standard model with a single central banker used to derive the Taylor rule does not predict interest rate autocorrelation. Since interest rates are autocorrelated in the data, the empirical analysis of Taylor rules usually involves the addition of lagged interest rates to the theoretical relation (see, for example, Clarida *et al.*, 1999).

¹⁹In order to get a more accurate picture of the distribution, these histograms were constructed using simulations of 10000 observations.

²⁰Because the transition matrix has a built-in inertia when the diagonal elements are non-zero and in order not to overstate the policy inertia predicted by the voting game, the histograms are plotted using only observations where there is a change in the state of nature.

²¹English *et al.* (2003) reports evidence that the autocorrelation in the U.S. Federal Funds Rate is the result of both policy inertia and shock persistence.

4 Comparing Monetary Policy Institutions

This paper shows the existence of a political failure in monetary policy making by committee. However, the fact that a fictional social planner can improve upon committees is no reason to conclude that this institutional arrangement is inefficient. In the real world, the only fair comparison is among political equilibria that can be obtained in the class of available institutions.²² In order to conclude that a given institution is inefficient, one must show that there exists another institution that increases the utilities of both policy makers. Unfortunately, this question cannot be answered in a definitive way because, for obvious reasons, the set of feasible institutions cannot be fully characterized.

In this section, we consider an alternative institutional arrangement that is identical to the one we have discussed so far, except for the fact that the default policy is fixed. This shuts down the dynamic link between periods and eliminates the rationale for not implementing the preferred policy in the even states of nature. Figure 7 shows the stationary policy rules when the default is either of the bliss points a through e . These policy rules do not depend on either the status quo or the matrix of transition probabilities. That is, when plotted as a function of q , they are horizontal lines and are the same for any transition matrix. The optimal proposal in the even states of nature is the preferred bliss point for any fixed default, but in the odd states it crucially depends on the location of the default. Table 1 reports the *ex-ante* (average) payoff for each committee member under three alternative institutions: 1) a committee with an evolving default; 2) a committee with a fixed default; and 3) full delegation to C .²³

Table 1 shows that a redistribution of utilities across members is obtained by varying the fixed default. Note that the “average” preferred policies in the states of disagreement for P and C are, respectively, b and d . This is why P (C) obtains a high payoff when b (d) is the fixed default. However, the best fixed default for the chairman is policy e . The reason is that this default always gives P a lower payoff than C 's ideal point. Thus, the chairman has enough bargaining power to propose his preferred point in all states of nature. Having a committee with a fixed default at e is therefore equivalent to an institution where C is the single central banker (see the last column in Table 1). Note that while having a fixed default eliminates the dynamic inefficiency in the states of agreement, it does not implement the efficient outcome in the states of disagreement. To see this, suppose, for example, that

²²For a discussion along the same lines, see Besley and Coate (1998, Section IV).

²³This comparison is meant to be suggestive only. The stylized nature of the model developed here prevents us from assessing more in detail the potential social welfare implications of committee decision-making, as well as the empirical relevance of dynamic inefficiency. We intend to take up these issues in future work.

the fixed default is c . Then, the committee selects c whenever $\varepsilon = \varepsilon_1$ or $\varepsilon = \varepsilon_5$ occur. This outcome is clearly inefficient because the utility of both committee members would increase by choosing a policy between the values preferred by P and C . In other words, there is inefficient risk sharing between P and C when the default is fixed. This source of inefficiency is also present in the model with an endogenous default, but it is less severe. The reason is that, when the default is endogenous, committee members can smooth their bargaining power across states of nature and, consequently, insure themselves against the eventuality of having little bargaining power in the next meeting.

Now, compare the average payoffs under committees with endogenous and fixed defaults. Clearly, an endogenous default lowers the average payoff to both members in the even states compared with a fixed default because the policy preferred by both members is not implemented. However, Table 1 shows that starting with an endogenous default, the committee would not agree on amending the institution because any choice of fixed default would lower the *ex-ante* utility of one of the policy makers. Table 1 also shows that, when the default is endogenous, C obtains a larger share of the surplus when the transition matrix is B rather than A . This is so because the optimal policy changes more drastically when shocks follow matrix B . Since the preferred policy in each period is more likely to be far from the previous policy, the chairman has more leverage in proposing his preferred instrument value.

Regarding the variance of policy decisions under these institutions, note that a committee with an endogenous status quo generally lowers the variance of both x and Δx compared with a single central banker and a committee with a fixed default (except when the default is c). This result is a consequence of the local policy inertia introduced by the endogenous status quo.

The absence of an institution that Pareto-dominates an arrangement with an endogenous default can explain its endurance, but it cannot explain why this institutional feature is observed so often in practice.²⁴ To answer this question, Riboni (2004) shows that, in a model without uncertainty, an endogenous default works as a commitment device and makes credibility problems less severe.

5 Summary

This paper models monetary policy making as a dynamic non-cooperative game. Committee members sequentially decide the policy for the period after observing the current realization of a preference shock. Depending on the shock, policy makers may agree or disagree about

²⁴On this point, see Tsebelis (2002, p. 8). Rasch (2000) identifies countries where an evolving default is part of the formal rules in legislative decision making.

the optimal monetary stance for the period. In this model, the first-best policy can be easily characterized: it satisfies a risk-sharing condition in the states of disagreement and prescribes the preferred policy of both agents in the states of agreement. This paper shows that, in the absence of commitment, committee decision making does not implement the first-best. Inefficiencies arise in all states of nature. In states of agreement, policy makers do not choose the policy they both currently prefer, because they face a trade-off between the benefit of selecting their preferred policy in the current period and the cost of reducing their bargaining power in the future. In states of disagreement, inefficiency is due to incomplete risk-sharing between committee members. Stochastic simulations show that committee decision making 1) induces policy smoothing in the sense that the policy variable changes less often than the state of nature and 2) endogenously generates autocorrelation in interest rates. Finally, we analyze committee decision making with a fixed default and show that this alternative arrangement removes the inefficiency in states of agreement by eliminating the incentive to smooth bargaining power across states. However, compared to a model with endogenous default, a fixed default model delivers more inefficient risk-sharing in the states of disagreement. This may be a probable reason why, despite the inefficiencies described, policy making in practice often features an evolving default.

Finally, we emphasize that this paper does not intend to play down the advantages of policy making by committees. We recognize that committee decision making has many desirable attributes. First, previous works show that committees can help overcome credibility problems. Sibert (2003) studies the conditions under which committees have more incentives to build reputation than do individual central bankers. In Dal Bó (2005), committee decision making under a supermajority voting rule is able to deliver an ideal balance between commitment and flexibility. Second, another body of literature sees information sharing as the main rationale for committee decision making. This argument goes back to the celebrated Condorcet jury theorem. For example, Gerlach-Kristen (2003) shows that in presence of uncertainty about potential output, voting by committees leads to more efficient signal extraction. Experimental studies by Blinder and Morgan (2000) and Lombardelli *et al.* (2005) provide some support for this conclusion.

**Table 1. Comparison of Voting Models
with Endogenous and Fixed Default**

Variable	Endogenous	Fixed					Single Banker
		<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	<i>e</i>	
<i>Matrix A</i>							
<i>C</i> 's mean payoff	-0.85	-0.59	-1.23	-0.92	-0.17	0	0
<i>P</i> 's mean payoff	-0.51	-0.92	-0.26	-0.59	-1.48	-1.99	-1.99
<i>Var</i> (<i>x</i>)	0.58	1.13	0.83	0.50	0.75	0.96	0.96
<i>Var</i> (Δx)	0.27	0.78	0.45	0.44	0.77	0.95	0.95
<i>Matrix B</i>							
<i>C</i> 's mean payoff	-0.68	-0.82	-1.00	-0.83	-0.17	0	0
<i>P</i> 's mean payoff	-1.05	-0.83	-0.33	-0.82	-1.49	-1.99	-1.99
<i>Var</i> (<i>x</i>)	0.75	1.29	0.90	0.58	0.75	0.96	.96
<i>Var</i> (Δx)	0.88	1.82	1.27	0.82	1.05	1.35	1.35

Notes: The numbers in this Table were computed using 10000 simulations.

A Committee with $n+1$ Members

Consider a committee composed of $n + 1$ members. Let n be odd and $[\underline{x}, \bar{x}]$ denote the policy space where policies take value. For a policy change, the chairman needs $(n + 1)/2$ favorable votes besides his own. Each member other than the chairman is indexed by j , with $j \in N = \{1, \dots, n\}$. When $\varepsilon = \varepsilon_1$, members disagree in their preferred instrument values, $r_j(\varepsilon_1)$. We order the n members other than the chairman so that member 1 (n) is the one with the smallest (largest) preferred value under shock ε_1 , and $r_1(\varepsilon_1) \leq r_2(\varepsilon_1) \leq \dots \leq r_n(\varepsilon_1)$. The median is the one with index $(n + 1)/2$. When $\varepsilon = \varepsilon_2$, all members agree and $b \in (\underline{x}, \bar{x})$ is their preferred point. We assume that c and a are, respectively, the preferred values of the chairman and the median, with $\underline{x} \leq a < b < c \leq \bar{x}$. As before, we assume that at time t the voting representative j accepts proposal x if and only if

$$U_j(x, \varepsilon_i) + \delta \sum_{k=1}^2 p_{ki} V_{j,t+1}(x, \varepsilon_k) \geq U_j(q, \varepsilon_i) + \delta \sum_{k=1}^2 p_{ki} V_{j,t+1}(q, \varepsilon_k).$$

This requirement, which is stricter than sequential rationality when $n \geq 3$, rules out equilibria where players accept a proposal they do not like for the simple reason that a single rejection does not affect the voting outcome (see Baron and Kalai, 1993). The acceptance set is then defined as

$$A_t(q, \varepsilon_i) = \{x \in [\underline{x}, \bar{x}] : |\{j \text{ accepts } x\}| \geq (n + 1)/2\}.$$

Denote by r_t^i a policy at time t under the shock ε_i . Note that the proposal made by the agenda setter concerns only the current period. However, in order to accept or reject the proposal, members implicitly compare two sequences of policies, where future policies are derived by using the proposal rules for subsequent periods.

Lemma 1. *Suppose $T \leq \infty$. Let $\{\tilde{r}_s^1, \tilde{r}_s^2\}_{s=t}^T$ and $\{\hat{r}_s^1, \hat{r}_s^2\}_{s=t}^T$ be two arbitrary policy sequences starting from an arbitrary t . The difference between the utilities associated with these two sequences is a monotone function of $r_j(\varepsilon_1)$.*

Proof: Without any loss of generality, suppose that the current shock is ε_1 . Write the expected utility associated with the sequence $\{\tilde{r}_s^1, \tilde{r}_s^2\}_{s=t}^T$,

$$E_t \left(\sum_{s=t}^T \delta^{s-t} U_j(\tilde{r}_s, \varepsilon_s) \right) = -(\tilde{r}_s^1 - r_j(\varepsilon_1))^2 - \delta p_{11} (\tilde{r}_{s+1}^1 - r_j(\varepsilon_1))^2 - \delta p_{21} (\tilde{r}_{s+1}^2 - b)^2 + \dots,$$

and with the alternative sequence $\{\hat{r}_s^1, \hat{r}_s^2\}_{s=t}^T$,

$$E_t \left(\sum_{s=t}^T \delta^{s-t} U_j(\hat{r}_s, \varepsilon_s) \right) = -(\hat{r}_s^1 - r_j(\varepsilon_1))^2 - \delta p_{11} (\hat{r}_{s+1}^1 - r_j(\varepsilon_1))^2 - \delta p_{21} (\hat{r}_{s+1}^2 - b)^2 + \dots$$

Compute the derivative of the difference of these two utilities with respect to $r_j(\varepsilon_1)$ and note that it does not depend on $r_j(\varepsilon_1)$. Then, the difference in utility among any two sequences is monotone in $r_j(\varepsilon_1)$. ■

From this lemma, it follows:

Result 1. *A proposal is accepted if and only if it is accepted by the median.*

Since the chairman only needs the approval of the median to pass a proposal and the preferences of the other members do not matter, then a committee with $n+1$ members is equivalent to a two-person committee with the chairman and the median as the only policy makers.

B Algorithm to Solve for Stationary Decision Rules

Step 1. Starting at time $t = T$, solve the chairman's optimization problem for a set of discrete nodes n_j , for $j = 1, 2, \dots, N$ in $[a, c]$, given the shock $\varepsilon = \varepsilon_i$, for $i = 1, 2$. The nodes n_j may be interpreted as possible status quo at the beginning of period T . Given n_j and ε_i , the chairman's problem at time $t = T$ is

$$V_{C,T}(n_j, \varepsilon_i) = \max_{x \in [a, c]} U_C(x, \varepsilon_i),$$

subject to the nonlinear constraint $U_P(x, \varepsilon_i) \geq U_P(n_j, \varepsilon_i)$. This maximization problem is solved numerically for each n_j and ε_i using a hill-climbing method. The result is a collection of $2N$ optimal proposal values $G_{C,T}(n_j, \varepsilon_i)$. Using these optimal values, compute $V_{C,T}(n_j, \varepsilon_i) = U_C(G_{C,T}(n_j, \varepsilon_i), \varepsilon_i)$ and $V_{P,T}(n_j, \varepsilon_i) = U_P(G_{C,T}(n_j, \varepsilon_i), \varepsilon_i)$ for all n_j and ε_i .

Step 2. For each ε_i , approximate the continuous value function $V_{C,T}(q, \varepsilon_i)$ using a Chebyshev polynomial of order $N - 1$. The polynomial coefficients are obtained from the Least Squares projection of $V_{C,T}(n_j, \varepsilon_i)$ on a constant and the first $N - 1$ members of the Chebyshev polynomial family. At the N nodes $q = n_j$, the Chebyshev polynomial fits $V_{C,T}(q, \varepsilon_i)$ exactly. For points $q \neq n_j$, the value of $V_{C,T}(q, \varepsilon_i)$ is computed by interpolation (*i.e.*, by evaluating the Chebyshev polynomial at q). For each ε_i , the value function $V_{P,T}(q, \varepsilon_i)$ is approximated likewise.

Step 3. Move backwards one period. For each possible status quo n_j and each possible shock realization ε_i , solve numerically the chairman's problem

$$V_{C,t}(n_j, \varepsilon_i) = \max_{x \in [a, c]} U_C(x, \varepsilon_i) + \delta \sum_{k=1}^2 p_{ki} V_{C,t+1}(x, \varepsilon_k),$$

subject to

$$U_P(x, \varepsilon_i) + \delta \sum_{k=1}^2 p_{ki} V_{P,t+1}(x, \varepsilon_k) \geq U_P(n_j, \varepsilon_i) + \delta \sum_{k=1}^2 p_{ki} V_{C,t+1}(n_j, \varepsilon_k),$$

where the value functions are replaced by their respective approximating polynomials. The result is a collection of $2N$ optimal proposal values $G_{C,t}(n_j, \varepsilon_i)$. Using these optimal values, compute $V_{C,t}(n_j, \varepsilon_i)$ and $V_{P,t}(n_j, \varepsilon_i)$ for all n_j and ε_i .

Step 4. Repeat Steps 2 and 3 backwards until the chairman's decision rules converge. ■

References

- [1] Alesina, Alberto, (1987), “Macroeconomic Policy in a Two-Party System as a Repeated Game,” *Quarterly Journal of Economics* 102, pp. 651-678.
- [2] Baron, David P., (1996), “A Dynamic Theory of Collective Goods Programs,” *American Political Science Review* 90, pp. 316-330.
- [3] Baron, David P. and Eehud Kalai, (1993), “The Simplest Equilibrium of a Majority Rule Game,” *Journal of Economic Theory* 61, pp. 290-301.
- [4] Baron, David P. and John Ferejohn, (1989), “Bargaining in Legislatures,” *American Political Science Review* 83, pp. 1181-1206.
- [5] Baron, David P. and Michael C. Herron, (2003). “A Dynamic Model of Multidimensional Collective Choice in Computational Models in Political Economy, Ken Kollman, John H. Miller, and Scott E. Page, eds. Cambridge, MA: The MIT Press.
- [6] Bernheim, Douglas, Antonio Rangel and Luis Rayo, (2005), “The Power of the Last Word in Legislative Policy Making.” *Econometrica*, forthcoming.
- [7] Besley, Timothy and Stephen Coate. (1998). “Sources of Inefficiency in a Representative Democracy: A Dynamic Analysis,” *American Economic Review* 88, 139-156.
- [8] Blinder, Alan S. and John Morgan, (2005), “Are Two Heads Better Than One? Monetary Policy by Committee,” *Journal of Money, Credit and Banking* 37, pp. 789-812.
- [9] Clarida, Richard, Jordi Gali and Mark Gertler, (1999), “The Science of Monetary Policy: A New Keynesian Perspective,” *Journal of Economic Literature* 37, pp. 1661-1707.
- [10] Chappell, Henry W., Rob Roy McGregor, and Todd Vermilyea, (2005), *Committee Decisions on Monetary Policy*, Cambridge: MIT Press.
- [11] Cukierman, Alex, (1991), “Why Does the Fed Smooth Interest Rates?” in Michael Belongia, ed., *Monetary Policy on the 75th Anniversary of the Federal Reserve System*, Boston: Kluwer Academic Publishers, pp. 111-47.
- [12] Dal Bó, Ernesto, (2005), “Committees With Supermajority Voting Yield Commitment With Flexibility,” *Journal of Public Economics*, forthcoming.

- [13] English, William, William Nelson, and Brian Sack, (2003), "Interpreting the Significance of the Lagged Interest Rate in Estimated Monetary Policy Rules," *Contributions to Macroeconomics* 3(1), Article 5.
- [14] Eijffinger, Sylvester, Eric Schaling, and Willem Verhagen, (1999), "A Theory of Interest Rate Stepping: Inflation Targeting in a Dynamic Menu Cost Model," Tilburg University, Mimeo.
- [15] Fry, Maxwell, DeAnne Julius, Lavan Mahadeva, Sandra Roger, and Gabriel Sterne, (2000), "Key Issues in the Choice of Monetary Policy Framework," in Lavan Mahadeva and Gabriel Sterne (eds), *Monetary Frameworks in a Global Context*, London: Routledge.
- [16] Gerlach-Kristen, Petra (2003). "Monetary Policy Committees and Interest Rate Setting," *European Economic Review*, forthcoming.
- [17] Goodfriend, Marvin, (1991), "Interest Rate Smoothing in the Conduct of Monetary Policy," *Carnegie-Rochester Conference Series on Public Policy* 37, pp. 7-30.
- [18] Hetzel, Robert L., (1998), "Arthur Burns and Inflation," Federal Reserve Bank of Richmond *Quarterly Review* 84, pp. 21-44.
- [19] Maisel, Sherman, (1973), *Managing the Dollar*, New York: W. W. Norton and Company.
- [20] Meyer, Laurence H., (2004), *A Term at the Fed*, New York: Harper Collins.
- [21] Lombardelli, Clare, James Proudman, and James Talbot, (2005), "Committee Versus Individuals: An Experimental Analysis of Monetary Policy Decision Making," *International Journal of Central Banking* 1, pp. 181-205.
- [22] Orphanides, Athanasios, (2003), "Monetary Policy Evaluation with Noisy Information," *Journal of Monetary Economics* 50, pp. 605-31.
- [23] Rasch, Bjorn Eric, (2000), "Parliamentary Floor Procedures and Agenda Setting in Europe," *Legislative Studies Quarterly* 25, pp. 3-23.
- [24] Riboni, Alessandro, (2004), "Committees as Substitute for Commitment," University of Montreal, Mimeo.
- [25] Romer, Thomas and Howard Rosenthal, (1978), "Political Resource Allocation, Controlled Agendas, and the Status Quo," *Public Choice* 33, pp. 27-44.

- [26] Sibert, Anne, (2003) “Monetary Policy Committees: Individual and Collective Reputations,” *Review of Economic Studies* 70, pp. 649-665.
- [27] Tsebelis, George, (2002), *Veto Players: How Political Institutions Work*, Princeton: Princeton University Press.
- [28] Judd, Kenneth L., (1998), *Numerical Methods in Economics*, Cambridge: MIT Press.
- [29] Waller, Christopher J., (2000), “Policy Boards and Policy Smoothing,” *Quarterly Journal of Economics* 115, pp. 305-339.
- [30] Woodford, Michael, (2003), “Optimal Interest-Rate Smoothing,” *Review of Economic Studies* 70, pp. 861-886.

Figure 1: Monetary Policy Decisions

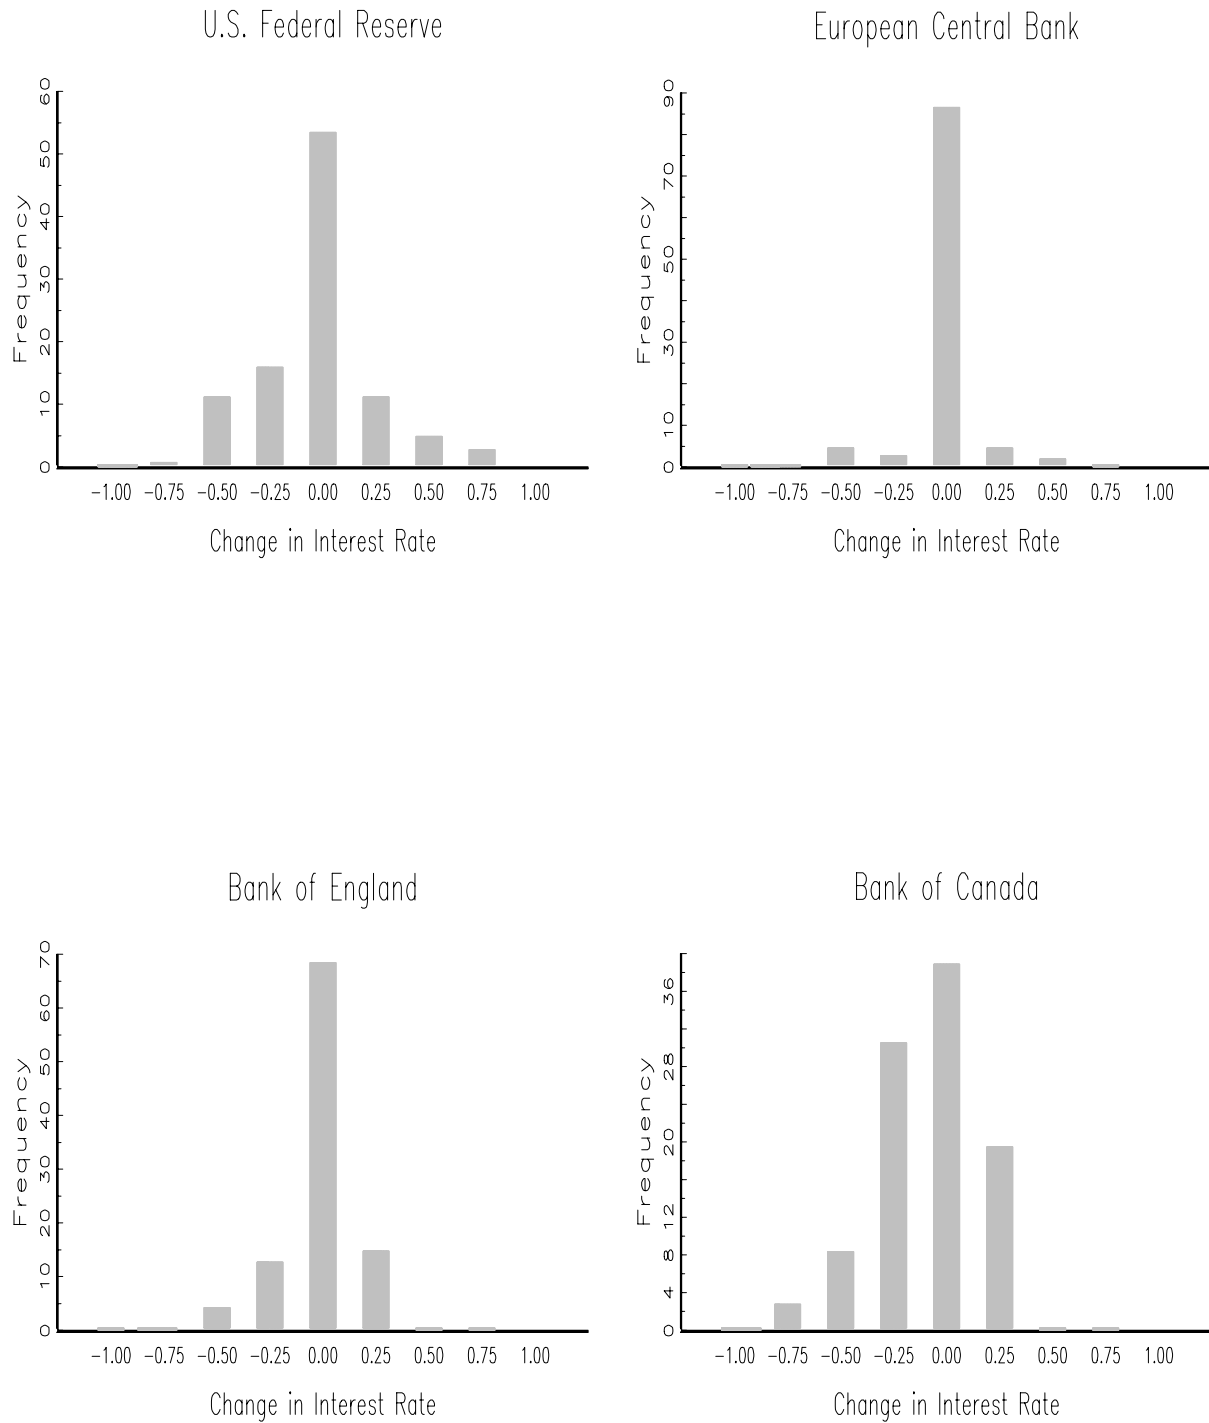


Figure 2: Policy Rules for Two-State Model

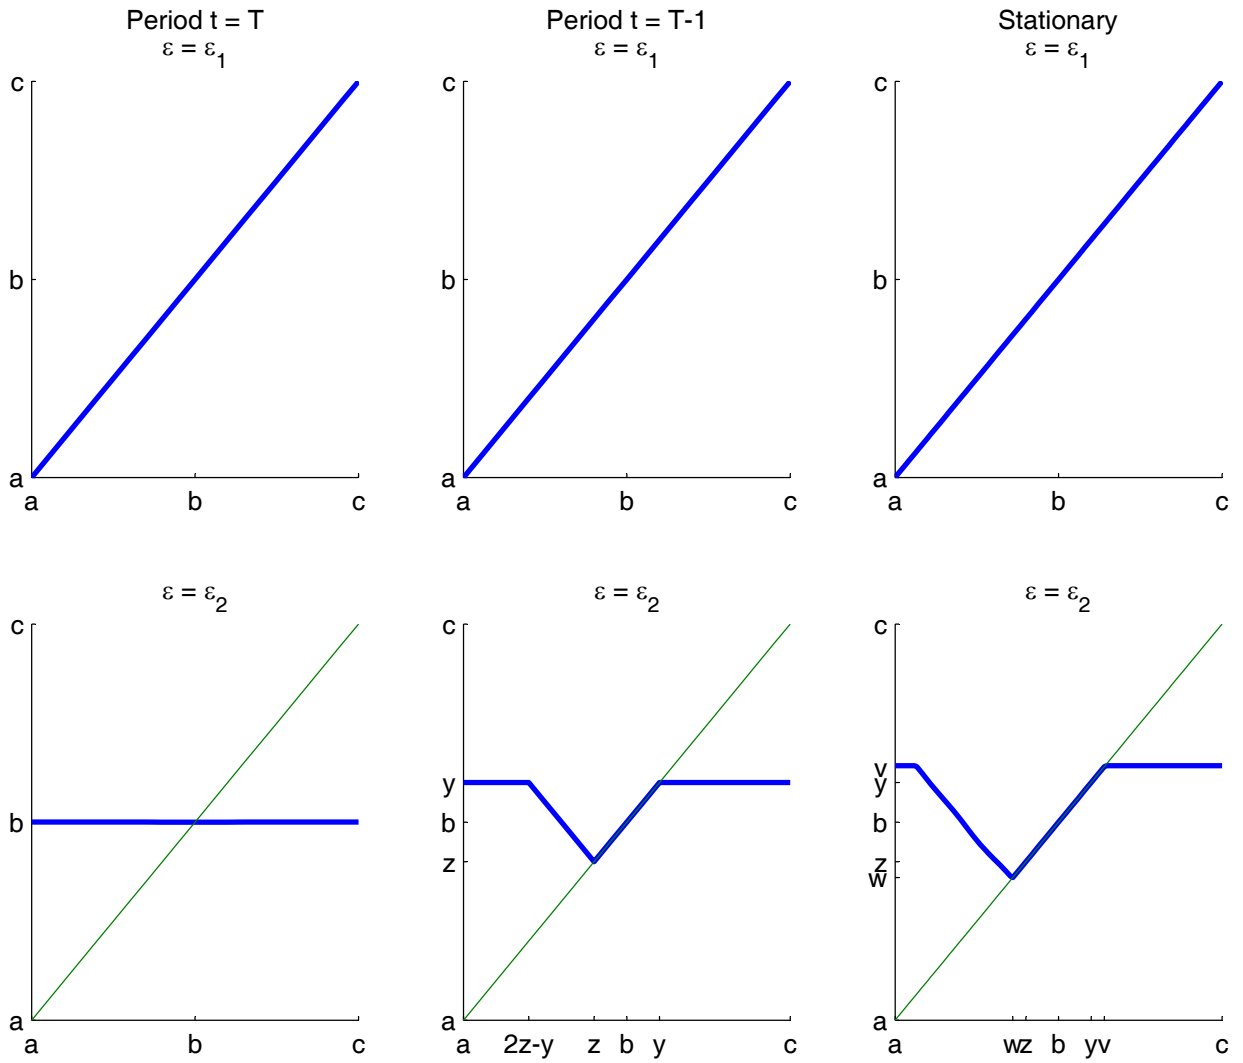


Figure 3: Stationary Policy Rules for Six-State Model

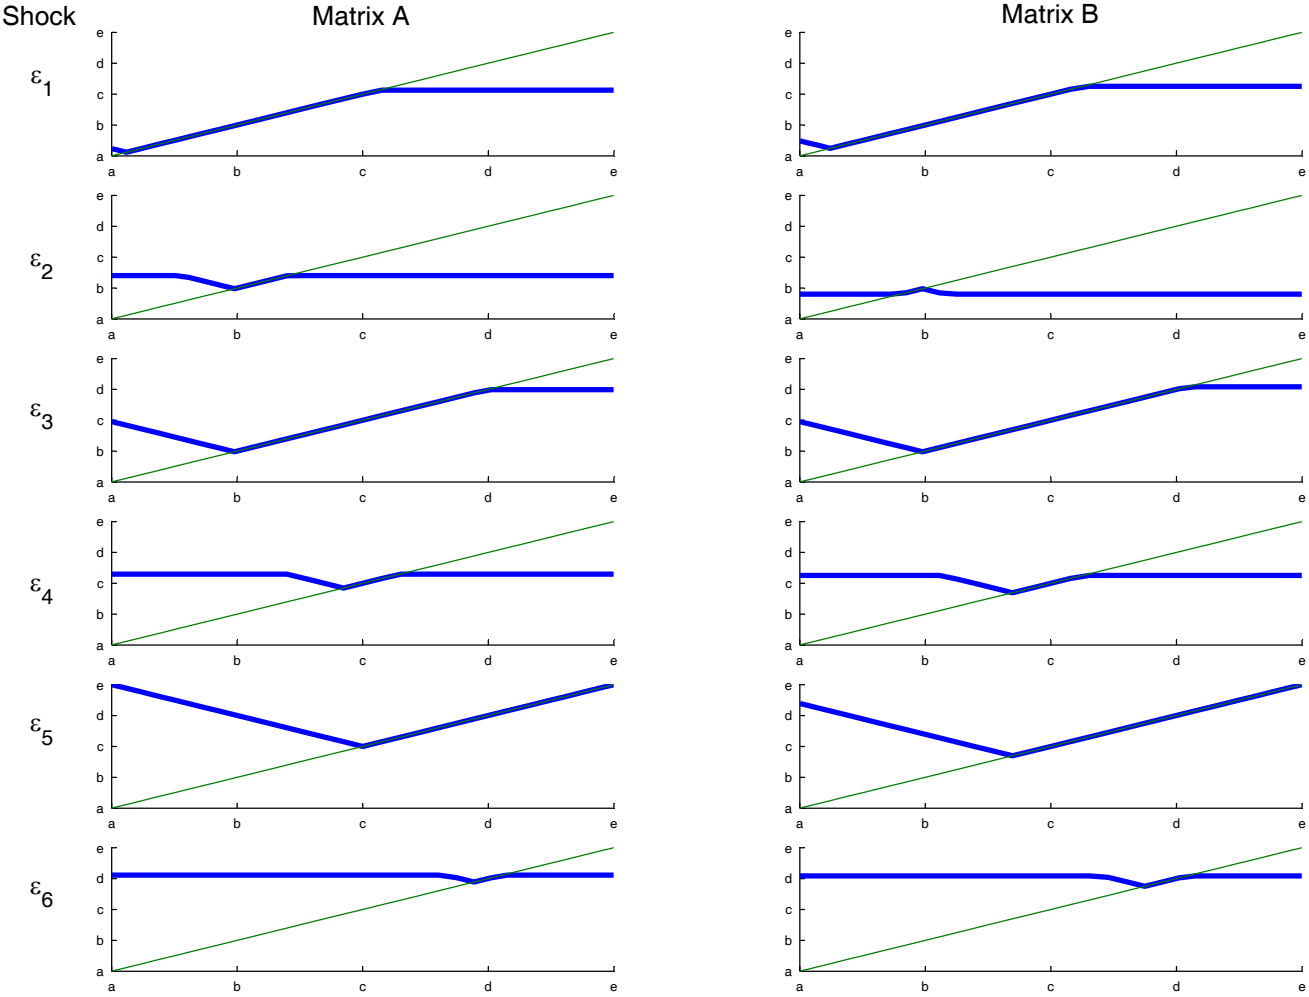


Figure 4: Simulations

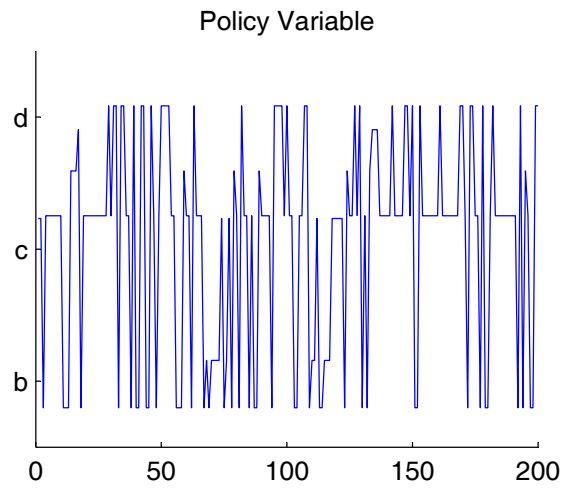
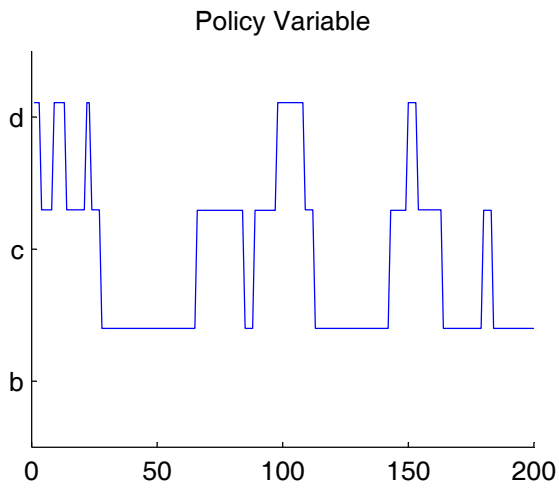
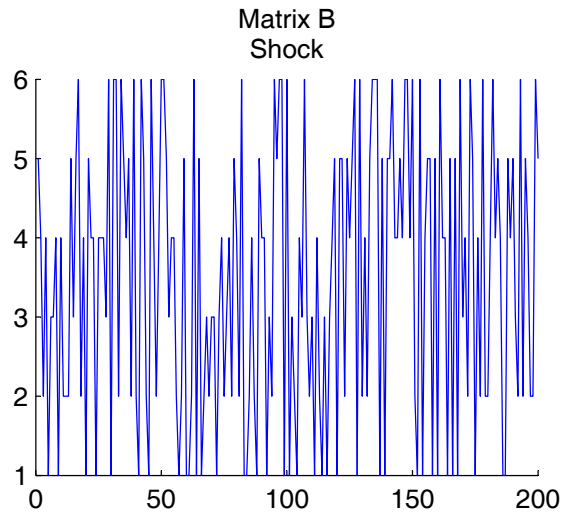
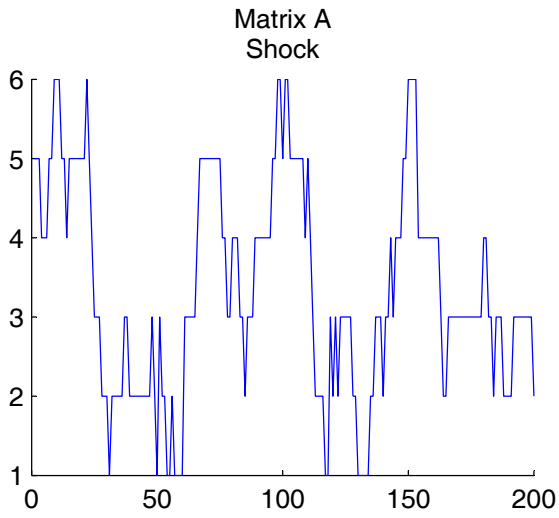


Figure 5: Histograms

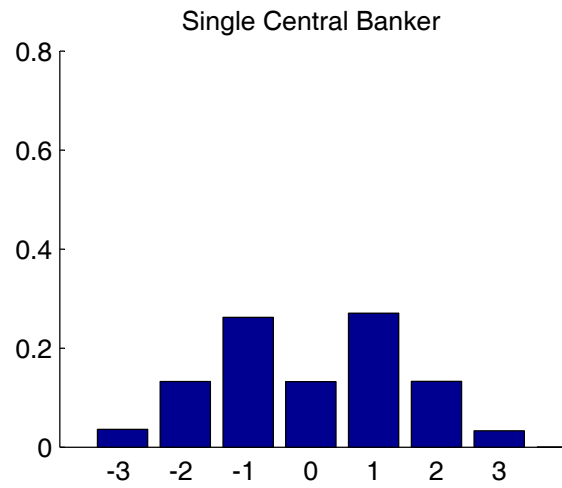
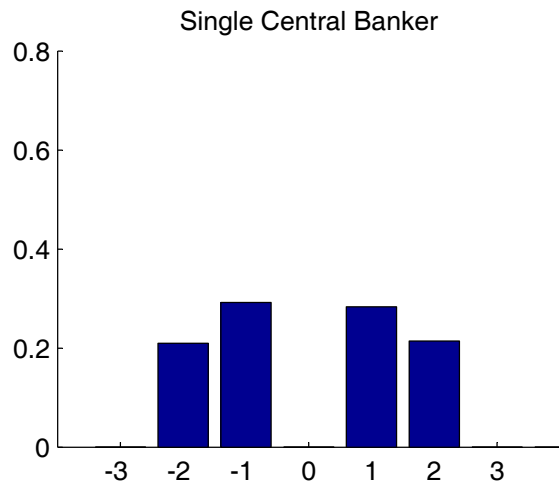
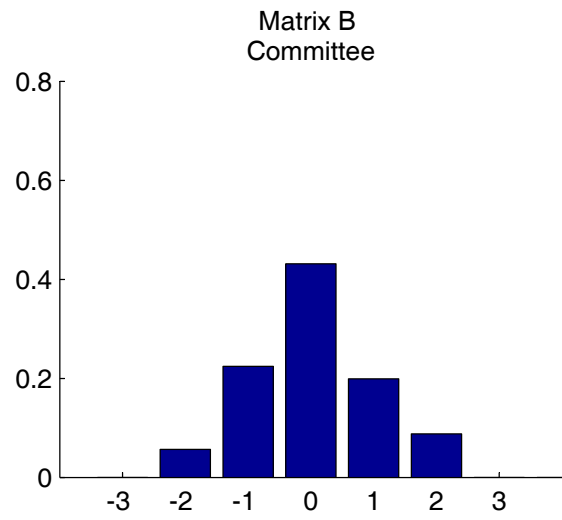
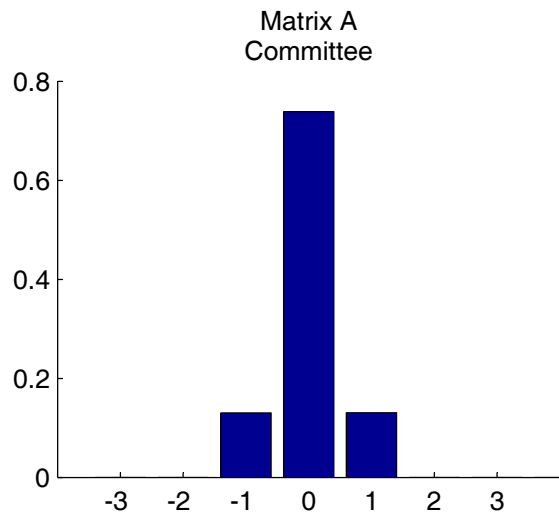


Figure 6: Autocorrelation Functions

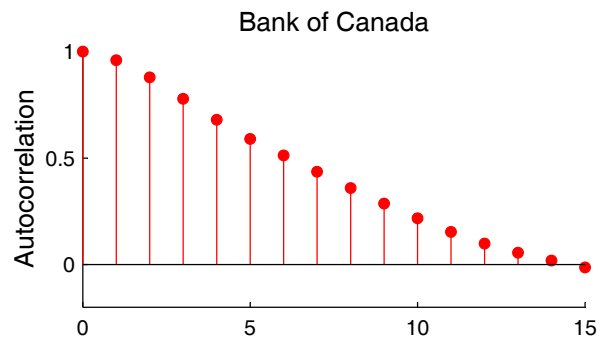
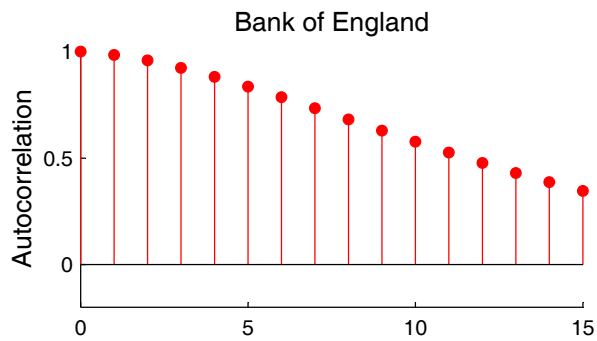
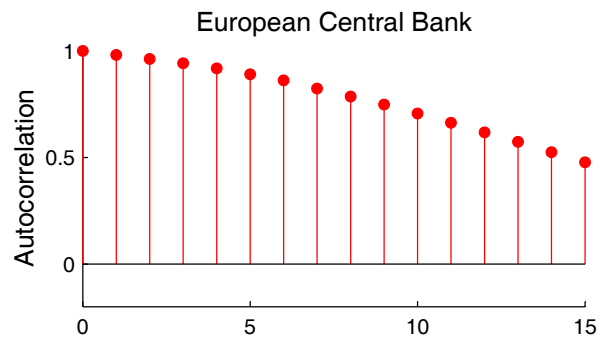
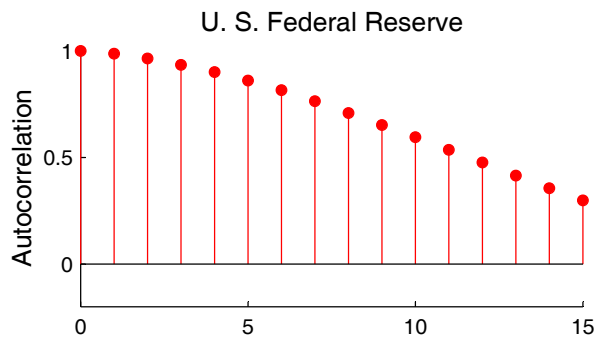
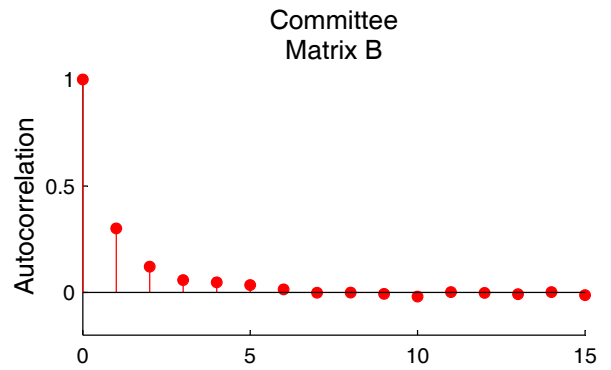
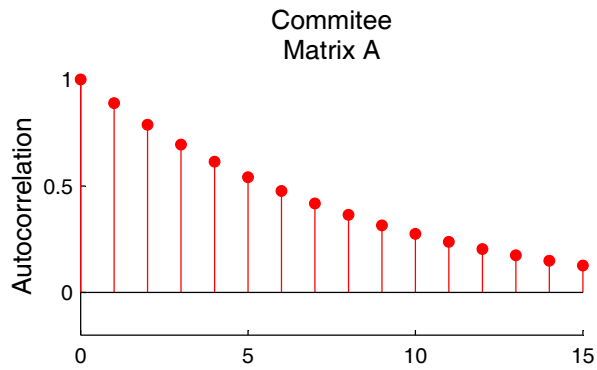


Figure 7: Stationary Policy Rules when Default is Fixed

