FEDERAL RESERVE BANK *of* ATLANTA

# Simple versus Optimal Rules as Guides to Policy

William A. Brock, Steven N. Durlauf, James M. Nason,
and Giacomo Rondina

Working Paper 2007-7
April 2007

WORKING PAPER SERIES

FEDERAL RESERVE BANK *of* ATLANTA | WORKING PAPER SERIES

# Simple versus Optimal Rules as Guides to Policy

William A. Brock, Steven N. Durlauf, James M. Nason, and Giacomo Rondina

Working Paper 2007-7
April 2007

**Abstract:** This paper contributes to the policy evaluation literature by developing new strategies to study alternative policy rules. We compare optimal rules to simple rules within canonical monetary policy models. In our context, an optimal rule represents the solution to an intertemporal optimization problem in which a loss function for the policymaker and an explicit model of the macroeconomy are specified. We define a simple rule to be a summary of the intuition policymakers and economists have about how a central bank should react to aggregate disturbances. The policy rules are evaluated under minimax and minimax regret criteria. These criteria force the policymaker to guard against a worst-case scenario, but in different ways. Minimax makes the worst possible model the benchmark for the policymaker, while minimax regret confronts the policymaker with uncertainty about the true model. Our results indicate that the case for a model-specific optimal rule can break down when uncertainty exists about which of several models is true. Further, we show that the assumption that the policymaker's loss function is known can obscure policy trade-offs that exist in the short, medium, and long run. Thus, policy evaluation is more difficult once it is recognized that model and preference uncertainty can interact.

JEL classification: E3, E5, E6

Key words: monetary policy rule, model uncertainty, minimax, minimax regret

---

The authors thank Bennett McCallum and Robert Tetlow for suggestions on an earlier draft. Brock and Durlauf thank the National Science Foundation for support. The views expressed here are the authors' and not necessarily those of the Federal Reserve Bank of Atlanta or the Federal Reserve System. Any remaining errors are the authors' responsibility.

Please address questions regarding content to William A. Brock, Madison Department of Economics, University of Wisconsin, 1180 Observatory Drive, Madison, WI 53706, wbrock@ssc.wisc.edu; Steven N. Durlauf, Madison Department of Economics, University of Wisconsin, 1180 Observatory Drive, Madison, WI 53706, sdurluaf@ssc.wisc.edu; James M. Nason, Research Department, Federal Reserve Bank of Atlanta, 1000 Peachtree Street, N.E., Atlanta, GA 30309, jim.nason@atl.frb.org; or Giacomo Rondina, Madison Department of Economics, University of Wisconsin, 1180 Observatory Drive, Madison, WI 53706, grondina@wisc.edu.

Federal Reserve Bank of Atlanta working papers, including revised versions, are available on the Atlanta Fed's Web site at www.frbatlanta.org. Click "Publications" and then "Working Papers." Use the WebScriber Service (at www.frbatlanta.org) to receive e-mail notifications about new papers.

## 1. Introduction

This paper contributes to the policy evaluation literature by developing new strategies to study alternative policy rules. We are interested in assessing the case that can be made for optimal rules as opposed to simple rules for policy. Our specific focus is on monetary policy rules. In our context, an optimal rule represents the solution to an intertemporal optimization problem in which a loss function for the policymaker and an explicit model of the macroeconomy are specified. We define a simple rule as a summary of the intuitions policymakers and economists have about how a central bank should react to aggregate disturbances.

Simple monetary policy rules have evolved since Friedman's classic (1948) presentation of the *k*% money growth rule. Modern versions of simple monetary policy rules typically embody a notion of "leaning against the wind".[2] A central bank plays an active role in stabilizing macroeconomic aggregates when it works to offset adverse shocks. The canonical modern example of a simple policy heuristic is Taylor's rule (1993), in which the Federal funds rate rises in response to output relative to trend and by more than one-for-one with the inflation rate. Taylor's rule is motivated by the notion that a central bank should attempt to reverse aggregate demand shocks, whereas aggregate supply shocks should be allowed to operate unhindered in the economy.

The distinction between optimal and simple rules has taken on particular significance in recent debates over the desirability of inflation targeting, as exemplified in the competing views expressed in Svensson (2003) and McCallum and Nelson (2005). The latter paper takes particular interest in the claim made by Svensson that targeting rules are preferred to instrument rules because the former embodies the restrictions intertemporal optimality places on the solution to the optimal policy problem.[3] As a

[2]Tobin (1983) distinguishes between fixed and reactive rules. Recent work on monetary policy evaluation focuses on the latter.

[3]Some arguments that appear in the literature strike us as self-evidently weak, for example, the claim that an instrument rule is inferior to a targeting rule because the latter uses the entire information set the policymaker has at hand. It is obvious that any information available for a targeting rule may be exploited by an instrument rule.

theoretical matter, McCallum has shown that in principle an instrument rule may arbitrarily well approximate a targeting rule, so long as the rules are conditioned on the same information set. Nonetheless, Svensson's claim has force when targeting rules are compared to some classes of simple rules (e.g., Taylor rules) in which we are interested.[4]

This paper questions the case for optimal rules at two levels. First, we consider the role model uncertainty plays in vitiating the superiority of optimal to simple rules. As argued by McCallum and Nelson (2005)[5], the optimality of a rule with respect to a given economic environment begs the question of its evaluation when the true model of the economy is not known. Our paper evaluates this issue by comparing model-specific optimal rules with simple alternatives. Second, we compare simple and optimal rules from the perspective of frequency-specific effects on uncertainty. Conventional policy assessments assume that the preferences of the policymaker are summarized by the variances of the aggregates under study. As we discuss below, such preferences ignore the disparate short-, medium- and long-run affects of policy. Since Brock and Durlauf (2004, 2005) and Brock, Durlauf, and Rondina (2007) show that the choice of a policy rule yields a frequency by frequency variance tradeoff, we argue that these tradeoffs should be reported to a policymaker, to account for different preferences with respect to fluctuations at different cycles. We believe that such comparisons contain useful information for policymakers.

Our analysis is in the spirit of Friedman (1948) who explicitly grounds his arguments for tying monetary policy to a fixed money growth rule on information limitations facing a central bank. Friedman concludes his analysis with

…I should like to emphasize the modest aim of the proposal. It does not claim to provide full employment…It does not claim to eliminate entirely cyclical fluctuations in output and employment. Its claim to serious consideration is that it provides a stable framework of fiscal and monetary action, that it largely eliminates the uncertainty…of discretionary actions by government authorities…It is not perhaps a proposal that one would consider at all optimum if our knowledge of the fundamental causes of cyclical fluctuations were

[4]Notice that the comparison of targeting rules to simple instrument rules renders moot the claim that targeting rules are more transparent than instrument rules.

[5]See McCallum (1988, 1999) for early discussions of model uncertainty and monetary policy that broadly study many of the questions we wish to consider.

considerably greater than I think them to be; it is a proposal that involves minimum reliance on uncertain and untested knowledge. (pg. 263)

Model uncertainty has, due to the pioneering work of Hansen and Sargent ((2003) and especially (2007)) on robustness, become a major feature of current research.[6] In this paper we follow an approach developed in Levin and Williams (2003) and Brock, Durlauf, and West (2003, 2006) to evaluate policies in the presence of model uncertainty. The idea of this approach is to first construct a model space that includes all candidate models for the economy, evaluate policies for each of the candidate models, and then determine how to draw policy influences given the fact that the true model is unknown. This approach differs from the usual robustness analysis in that we do not focus on model misspecification which is measured local to a given baseline model. Rather, our approach acknowledges the global character of model uncertainty.

Given a model space, it is necessary to determine how to aggregate information on rule performance across the different models. In this paper, we consider non-Bayesian approaches that are based on minimax and minimax regret criteria. These approaches involve guarding against “worst case scenarios”, but do so in different ways that we formalize below. Relative to the Bayesian approach described below, neither minimax nor minimax regret requires the policymaker to take a stance on model space priors in order to compute posterior model probabilities.

Our analysis *does not* address the question of the design of optimal rules in the presence of model uncertainty. We regard this as an important, but distinct question from understanding the sensitivity of rules to model uncertainty. We believe our exercise is useful because it provides insights into how to evaluate the performance of rules in the presence of model uncertainty that cannot be explicitly described in advance. There are fundamental conceptual questions involved in thinking about policy effects when one cannot specify the support of the uncertainty faced by a policymaker. Nevertheless, we believe insights into this difficult problem are found from understanding how rules perform within a model space when, by construction, the rules fail to account for model uncertainty. Our approach to non-Bayesian decision theory follows Hansen and Sargent

[6]Examples that focus on monetary policy include Giannoni (2002), Onatski and Stock (2002) and Tetlow and von zur Muehlen (2001).

(who focus on minimax evaluation). Since we operate within somewhat different model spaces, it requires additional justification for our approach.

In terms of substantive conclusions, we find that for some standard univariate macroeconomic models there exist significant limitations to model-specific optimal rules. When model uncertainty is restricted to the case where the target variable (inflation in our example) is either entirely forward-looking or backwards-looking, minimax and minimax regret often choose rules that do not feedback on the lagged instrument variable (i.e., the output gap) when the policymaker aims to minimize the total unconditional variance of the target. In contrast, the optimal rule for the forward looking model is favored when attention is restricted to the higher frequencies, even when the true model is uncertain. When we consider a hybrid model space, in which the relative weights of the backwards- and forward-looking components are unknown, the comparison of a model-specific optimal rule (based on the model where each component has equal weight) and a rule of thumb reveals that the model-specific rule is preferred when one focuses on overall variance and higher frequencies, whereas the rule of thumb is preferred for lower frequencies. Taken as a whole, these results indicate how the case for a model-specific optimal rule can breakdown when one relaxes the assumption that the true model is known as well as the assumption that the appropriate loss function is known.

Section 2 describes our basic framework for assessing model uncertainty. Section 3 discusses the evaluation of frequency-specific effects of policies. Section 4 applies our general ideas to policy evaluation when there is uncertainty about the respective importance of backwards- and forward-looking elements to the inflation process. Section 5 provides summary and discussion of future research directions. A Technical Appendix which provides derivations is available from the authors upon request.

## 2. Model uncertainty: basic ideas

At an abstract level, the issue of model uncertainty is easily described. Suppose that a policymaker faces a choice among a set $P$ of candidate policies. Given $\Theta$, which captures whatever features of the economy matter to the policymaker, a policymaker experiences a loss associated with those features

$$l(\Theta). \tag{1}$$

The features defined by $\Theta$ are not, for interesting cases, a deterministic function of a policy and so are in general described by a conditional probability

$$\mu\left(\Theta \middle| p, d, m\right) \tag{2}$$

where we explicitly allow the state of the economy to depend on the policy choice, $p$, the available data, *d*, and the model of the economy that is assumed to apply, *m*. Data dependence in this conditional probability may incorporate both direct dependence, as occurs in an autoregression, where forecasts depend on past realizations, or uncertainty about parameters. For simplicity, we measure losses as positive, so a policymaker wishes to minimize (1) by choosing a policy to affect (2). Our goal in the analysis is to relax the assumption that the policymaker knows the correct model of the economy, *m*. Rather, we assume that the policymaker knows that the correct model lies in a model space *M*. This space reflects uncertainty about the appropriate theoretical and functional form commitments needed to allow quantitative analysis of the effects of a policy.

In this paper we focus on non-Bayesian decision criteria with respect to model uncertainty.[7] In order to understand why we take this route, recall that the standard Bayesian solution to the optimal policy problem solves

$$\min_{p\in P} E\left(l(\Theta)\middle| p, d, M\right) = \min_{p\in P} \sum_{m\in M} E\left(l(\Theta)\middle| p, d, m\right)\mu\left(m \middle| d\right) \tag{3}$$

[7]Hansen and Sargent (2007) provide an extensive discussion on non-Bayesian approaches to macroeconomic analyses of the type we explore. Their analysis develops a comprehensive theory of robust decisionmaking. See also Brock, Durlauf and West (2003), whose formulation of model uncertainty and rule comparison we follow here.

where $\mu\left(m|d\right)$ denotes the posterior probability that model $m$ is the correct one.[8,9] Recalling that

$$\mu\left(d|m\right) \propto \mu\left(m|d\right)\mu\left(m\right) \tag{4}$$

the Bayesian approach thus requires explicit assignment of model priors. The assignment of these priors is problematic since a researcher rarely possesses any meaningful prior knowledge from which priors may be constructed. The default prior in the model uncertainty literature assigns equal probabilities to all elements of the model space; but as discussed in Brock, Durlauf, and West (2003) such an approach is hard to justify. Hence, a first limitation of Bayesian approaches to model uncertainty is the requirement of introducing priors that have little meaning. There has been work on evaluating the robustness of Bayesian inferences to prior choice in the statistics literature as well as efforts in the econometrics literature to develop alternatives to the diffuse priors that are used in the model uncertainty literature, but neither of these approaches is well understood for the context we study. We should note that the problem of defining priors continues to be a barrier to the adoption of Bayesian methods.[10]

Second, beyond the difficulty of specifying priors, we feel that there are other good reasons, for at least some exercises, to avoid weighting models by posterior weights. These have to do with the interpretation of the likelihood component in (4), i.e. $\mu\left(m|d\right)$. In our judgment, the utility of models is context dependent. A model which fits the overall data well, i.e. receives a high value of $\mu\left(m|d\right)$, may perform relatively poorly when employed to consider regime changes when compared to another model whose overall fit is poorer. It is easy to see how this could be the case when the model which fits historical data less well is relatively immune to the Lucas critique.

---

[8]We do not address the issue of how to interpret these probabilities when none of the models is true. This is a deep and unresolved problem; see Key, Perrichi and Smith (1998) for some efforts to address.

[9]The Bayesian solution to model uncertainty involves model averaging, an idea suggested in Leamer (1978) and developed in detail in Raftery, Madigan, and Hoeting (1997).

[10]Freedman (1991) has rejected the Bayesian approach for this reason.

Third, it is important to recognize that the model spaces studied in contemporaneous macroeconomics have not arisen *sui generis*; they reflect a history of theoretical and data work and as such, their relative fits reflect the interaction of macro modeling and empirical work. Nonetheless, we study a class of macro models that evolved to achieve better fit because of data mining. More generally, one should not reify a model space as representing the unique set of candidate models that were available to a researcher prior to data analysis. Instead, it should be recognized that the econometric versions of macro models that are under consideration (at a moment in time) have been influenced by the path of empirical work.

Finally, there are reasons to believe that individual preferences may treat model uncertainty differently from other types of uncertainty. These arguments are often based on experimental evidence of the type associated with the Ellsberg paradox, in which individuals appear to evaluate model uncertainty differently from other types. Put differently, experimental evidence suggests that preferences exhibit ambiguity aversion as well as risk aversion, when model uncertainty is considered. For our purposes, what matters is that the evidence of ambiguity aversion suggests a tendency to guard against the worst case scenario with respect to models; this idea is developed in Hansen and Sargent (2007) and Brock, Durlauf and West (2003).

Non-Bayesian decision criteria attempt to avoid dependence of inferences on model probabilities guarding against especially bad outcomes. At an intuitive level, these approaches abandon the search for optimal rules (in the Bayesian decision theory sense) in favor of good rules, where good is equated with the notion that the rule works relatively well regardless of which model is true. The best known approach is the minimax criterion, proposed by Wald (1950), under which a policy is chosen to minimize the expected loss under the least favorable model (relative to that rule) in the model set. The minimax policy choice is, for our formulation

$$\min_{p \in P} \max_{m \in M} E\left(l\left(\Theta\right) \middle| p, d, m\right). \tag{5}$$

The minimax approach has been pioneered by Hansen and Sargent for macroeconomic contexts and forms the basis of the literature on robustness analysis. From the perspective of the economic theory literature, versions of minimax have been justified by Gilboa and Schmeidler (1989) and Epstein and Wang (1994).

A standard criticism of the minimax criterion is that it is inappropriately conservative because it assumes the worst case possible in assessing policies. The force of this criticism, in our view, depends on whether the policy choice is driven by a relatively implausible model. This of course leads back to the question of model probabilities, the reasons for eschewing we have already given. Nevertheless, the extreme conservatism of minimax leads us to consider an alternative criterion: minimax regret.

Minimax regret was originally proposed by Savage (1951) to avoid the pessimism of minimax by moving away from the idea that the worst outcome should be the benchmark comparison. Rather, the focus in minimax regret is on the relative loss associated with choosing a policy in absence of knowledge of the correct model. For a given policy *p* and model *m*, the regret $R(p,d,m)$ associated with the policy choice is defined as

$$R(p,d,m) = E\left(l(\Theta)\middle| p,d,m\right) - \min_{p\in P} E\left(l(\Theta)\middle| p,d,m\right). \tag{6}$$

Regret therefore measure the loss incurred by a policy given a model relative to what would have incurred under the optimal policy for that model. The minimax regret policy is correspondingly defined by

$$\min_{p\in P} \max_{m\in M} R(p,d,m). \tag{7}$$

Seminal work by Manski (2005) has generated a great deal of attention in the microeconometrics literature for minimax regret.[11] As far as we know, minimax regret

[11]Chamberlain (2001) is an important precursor in suggesting minimax regret. See Brock (2006) and Manski (2006a,b) for applications.

has not been used to study macroeconomic questions in general, let alone with respect to questions associated with policy evaluation.[12]

The minimax regret criterion's appeal for policy comparison may be seen in the following example. Suppose that there are two policies and two possible models. For policy 1, there is a loss under model 1 of 22 and a loss under model 2 of 21. For policy 2, the losses under models 1 and 2 are 23 and 3 respectively. Under the minimax criterion, policy 1 is chosen since the maximum loss under it is 22 versus a maximum loss of 23 under policy 2. Under minimax regret policy 2 is chosen, since the maximum regret for policy 1 is 18 whereas for policy 2 it is 1. The reason why minimax regret chooses a different model than minimax is that under model 1, both policies lead to similar losses, whereas under model 2, the losses are quite different.

This example illustrates a key source of the appeal of minimax regret, namely, the criterion "normalizes" policy comparisons to account for the fact that certain realizations of the unknown state of nature (in our case, the correct model) are associated with relatively high losses regardless of the policy choice. In the language of classical decision theory, minimax regret has the virtue that it subtracts off for the loss associated with each candidate action the inevitable loss associated with a particular model, i.e. the loss under the model-specific optimal rule.

One objection to the minimax regret criterion is that decisionmaking under it is not required to obey the axiom of independence of irrelevant alternatives (IIA), something first shown in Chernoff (1954). It is not clear that this axiom is appropriate for policy decisions, as opposed to individual ones. Without focusing on the individual/policymaker distinction, we can see reasons why the axiom may not be natural. The axiom in essence states that one's preferences over different pairs of actions does not depend on the context in which they are made, where the total choice set is part of the context. This is not an obvious requirement of rationality.[13] In absence of the

[12]After writing this paper, we discovered Eozenou, Rivas and Schlag (2006) who provide a brief application of minimax regret to a question in economic growth to illustrate methods the authors have developed and as such is a macroeconomic example.

[13]Blume, Easley, and Halperin (2006) argue, for example, that an appropriate axiomatization of rationality should not require that agents have preferences which obey certain consistency conditions over all possible states of the world, as is required by

axiom, there do exist coherent axiom systems that produce minimax regret as the standard for evaluation; Stoye (2006) is the state of the art treatment. We note that in Stoye's system, for example, minimax regret obeys a weakened version of IIA called independence of nonoptimal alternatives, which means that the introduction or removal of a possible choice that is never the optimal one for any state of nature, will not matter in comparing the other choices.

We do not predicate our interest in minimax regret as a method of comparing policies on a strong adherence to any particular argument against the IIA assumption. We report minimax comparisons in addition to minimax regret comparisons to see where the criteria differ in practice. As will be apparent, minimax regret and minimax produced qualitatively similar results. We do not interpret this as indicating that one of the criteria is redundant, rather we feel it reinforces our general conclusions.

## 3. Frequency-specific effects of policies

In this section we consider rule performance with respect to fluctuations at different frequencies. Standard analyses of the effects of policies focus on weighted averages of the unconditional variances of the states and possibly the controls. Hence, it is common to compare policies by considering, for example, a weighted average of the unconditional variances of inflation and unemployment. Such calculations thus mask the effects of policies on the variance of fluctuations at different frequencies.

Brock and Durlauf (2004, 2005) and Brock, Durlauf, and Rondina (2007) argue that fundamental tradeoffs in the variance associated with different frequencies of state variables can arise for macroeconomic contexts of the type conventionally studied. We describe these tradeoffs for a one dimensional system

$$x_t = \beta E_t x_{t+1} + A(L)x_t + B(L)u_t + W(L)\varepsilon_t, \tag{8}$$

Savage, but rather that rationality requires certain consistency conditions only over states under consideration. While they do not explicitly pursue this approach as a critique of irrelevance of independent alternatives, their work can be used to justify its exclusion as a rationality requirement.

with feedback rule

$$u_t = \pi(L) x_{t-1}. \tag{9}$$

Let $f_x^C(\omega)$ denote the spectral density of the state when control rule $C$ is applied. Let $x_t^{NC}$ denote the system when $u_t \equiv 0$, i.e. there is no control; $f_x^{NC}(\omega)$ is the associated spectral density matrix.[14] Each choice of control can be understood as producing a sensitivity function $S(\omega)$ such that

$$f_x^C(\omega) = |S(\omega)|^2 f_x^{NC}(\omega). \tag{10}$$

Different choices of the feedback rule thus shape the spectral density of the state variables relative to the no control baseline. Hence, one can characterize what feedback rules can achieve in terms of frequency-by-frequency effects on the state by characterizing the set of sensitivity functions that are available to the policymaker. This is the key notion underlying the study of design limits for policy choice.[15]

It turns out that there is a simple way to characterize the set of feasible sensitivity functions in univariate systems. For each specification of a model of the uncontrolled system *M* and a control *C*, it must be the case that the associated sensitivity function obeys

$$\int_{-\pi}^{\pi} \log\left(|S(\omega)|^2\right) d\omega = \kappa(M, C). \tag{11}$$

[14]When the uncontrolled system contains explosive or unit roots in the AR component, then the spectral density will not exist. Our results on restrictions on the sensitivity function still hold in this case.

[15]Brock and Durlauf (2004, 2005) give initial treatments of these issues. Brock, Durlauf, and Rondina (2007) provides a detailed theory.

As discussed in Brock, Durlauf, and Rondina (2007)[16], if the system has no forward-looking elements, i.e. $\beta \equiv 0$, then the constraint $\kappa(M,C)$ depends only on the unstable autoregressive roots of the unconstrained system. The constraint equals 0 if there are no such roots, and is positive otherwise. The constraint is thus independent of the choice of control. Since $\log\left(|S(\omega)|^2\right)$ cannot be less than 0 for all frequencies, it must be the case that $\left|S(\omega)\right|^2$ cannot be less than 1 for all frequencies. Therefore, it is impossible to induce uniform reductions in the spectral density of the state variable relative to its value when there is no control, i.e. reductions at some frequencies require increases at others. Of course, the overall variance may be reduced.

The results are more complicated for forward-looking systems. When $\beta \in (0,1)$, it is possible for a rule to reduce the variance of the state at all frequencies, although whether this is possible will depend on the values of the polynomials in (8). Brock, Durlauf and Rondina (2007) show that a variance-minimizing rule may increase variance at some frequencies, even when a uniform reduction at all frequencies is possible.

Should a central bank evaluate rules on the basis of frequency specific effects? One reason why a central bank should do so concerns problems of measuring low frequency components of state variables, which is discussed in Onatski and Williams (2003).[17] Low frequency behavior is of course hard to identify. Such measurement problems take on particular importance if it is believed that long-run movements in the data are those outside the control of the central bank. For these reasons, it may be reasonable for a policymaker to compare the behavior of rules over integrals of "trimmed spectral densities," i.e. integrating spectral densities over $[-\pi,-\varsigma]\cup[\varsigma,\pi]$ instead of $[-\pi,\pi]$.[18]

[16]The results in Brock, Durlauf and Rondina (2007) for backwards-looking systems modulo details, may be found in the control theory literature; the results on forward-looking systems are new.

[17]Brock, Durlauf, Nason, and Rondina (2007) pursue this issue with respect to the effects on policy evaluation of alternative detrending methods.

[18]The issue concerns integrating spectral densities to compute losses, not to detrend data. In which case, the argument does not justify use of ad hoc filters.

A second argument for considering frequency-specific effects is based on nonseparable preferences. Under nonseparable preferences, a policymaker assigns different weights to different frequencies to calculate expected welfare. Otrok (2001) has a useful example where the weights on different frequencies can differ by more than 9:1. According to Otrok, if preferences exhibit habit persistence the loss associated with volatility at high frequencies is greater than the loss associated with volatility at low frequencies.

Otrok's results suggest that the study of frequency-specific losses may be important in understanding the role of preference uncertainty in evaluating policy rules. We can imagine two levels where such uncertainty exists. First, an analyst may not be certain about the preferences of a policymaker; hence he will want to report frequency-specific effects so that his findings may be used by policymakers with different preferences. Second, a policymaker who is uncertain about public preferences over, say, inflation and the output gap, may wish to account for this uncertainty in evaluating policy, i.e. to account for lack of knowledge of whether or not these preferences are separable. In this case, a policymaker will want to understand how losses occur for a policy at different frequencies.

The upshot of this discussion is that we believe part of the communication exercise for policymakers should involve reporting frequency specific comparisons.

## 4. Model uncertainty and monetary policy evaluation: some illustrations

In this section, we consider two analyses of rule of thumb versus optimal policies where model uncertainty is present. We follow Levin and Williams (2003) in placing primary focus on model uncertainty with respect to backwards- and forward-looking expectations in Keynesian models. In the context of the Phillips curve, the 2005 special issue, "*The econometrics of the New Keynesian price equation*", in the *Journal of Monetary Economics* reveals little consensus on this issue. For example, the Rudd and Whelan paper is especially clear about how disagreements over inflation dynamics create uncertainty for new Keynesian model building. The difference between backwards- and forward-looking models represents the modern instantiation of disagreements over the nature of the Phillips curve that helped launch the rational expectations revolution.

In our analysis, we consider a univariate system where inflation rate $\pi_t$ is controlled by the output gap $y_t$. Model-specific optimal rules will be constructed on the basis of assuming a particular law of motion for the inflation rate and the intertemporal loss function

$$E_0\left[\frac{1}{2}\sum_{t=0}^{\infty}\delta^t\pi_t\right]. \tag{12}$$

This corresponds to an "inflation nutter" in that volatility in the control variable is not considered. In subsequent work we consider richer loss functions, notably ones where the volatility of the control also matters, but this one is useful in illustrating our main ideas.

**a. theory uncertainty: backwards- versus forward- looking models**

We first consider the evaluation of policies when model uncertainty is represented by alternative forward-looking and backwards-looking theories of inflation. Our model space thus contains two candidates. The first is a backwards-looking model

$$\pi_t = a\pi_{t-1} + b_B y_{t-1} + \varepsilon_t;\ \varepsilon_t \sim WN. \tag{13}$$

The optimal feedback rule for this model $O_B$ is shown in the Technical Appendix[19] to be

$$y_t = -\frac{a}{b_B}\pi_t. \tag{14}$$

[19] The optimal policy calculations performed in this section and Section 4.c are similar to those found in Clarida, Gali, and Gertler (1999), especially for forward-looking models.

This is the same as the optimal feedback rule that minimizes the unconditional variance of the state. Notice that when the disturbance is not white noise the optimal rule will contain lagged state variables.

Our alternative forward-looking model is

$$\pi_t = \beta E_t \pi_{t+1} + b_F y_{t-1} + \varepsilon_t;\ \varepsilon_t \sim WN. \tag{15}$$

The optimal feedback rule ( $O_F$ ) is shown in the Technical Appendix to be

$$y_t = -\frac{\beta}{\delta b_F}\pi_t + \frac{\beta}{\delta} y_{t-1}. \tag{16}$$

Eq. (16) reflects the substantial smoothing typical of timeless perspective optimal rules in forward looking models. An explicit expression for $y_t$ in terms of current and past $\pi_t$ can be derived if $\beta < \delta$ .

In addition to $O_F$ and $O_B$ we consider two additional rules. The first is a rule of thumb ( $RoT$ )

$$y_t = -\bar{f}\pi_t \tag{17}$$

where, in addition to the negative sign, our choice of the coefficient $\bar{f}$ is such that the policy instrument is reacting to innovations in the inflation series more than one for one.[20] Second, we consider a restricted optimal rule ( $RO_F$ ) for the forward-looking model, by which we mean rules restricted to the form $y_t = f\pi_t$ in which the parameter $f$

[20]In this sense there is an analogy with the celebrated “Taylor principle” for monetary policy rules. The Taylor principle is defined in contexts where the policy rule is a nominal interest rate rule and the interest rate responds more than one for one to inflation. This ensures that when inflation increases the *real* interest rate increases because of the policy reaction. Thus, our focus on a single-equation model for inflation where the output gap is the policy instrument makes the analogy with the Taylor principle a loose one. The key property for our rule of thumb is that the instrument “leans against the wind” and that this reaction is relatively strong.

is chosen to minimize (12), which follows a candidate rule initially suggested in Levin and Williams (2003). The restricted optimal rule is shown in Brock, Durlauf and Rondina (2007) to take on a value of the coefficient $f$ so that

$$y_t = \frac{\left(1-\sqrt{1+8\beta^2}-4\beta^2\right)}{b_F 8\beta^2}\pi_t . \tag{18}$$

We now consider how these different rules perform given uncertainty about whether the true model is backwards-looking or forward-looking. We focus on the parameterizations $\frac{\beta}{\delta}=0.9$, $a=0.9$, and $b_B=b_F=.1$ (so that $\frac{b_B}{b_F}=1$) for the state equation and employ the normalization $\sigma_\varepsilon^2=1$.[21]

We study a rule of thumb motivated by uncertainty about the backwards- versus forward-looking specifications (13) and (15) rather than uncertainty about the parameters $b_B$ and $b_F$ within the models. This suggests the employment of a rule of thumb ($RoT$) that is normalized with respect to the slope parameter $b$ where $b=(1-\mu)b_B+\mu b_F$, for $\mu\in[0,1]$. We choose $y_t=\frac{-1.25}{b}\pi_t$;[22] the magnitude of the feedback parameter ensures that the reaction of the control to the state is relatively strong.

---

[21]Our choices of $b_F$ and $b_B$ are motivated by Woodford's (2003) preferred estimates for the slope of the New-Keynesian Phillips Curve of 0.096 and the Rudebusch and Svensson (1999) preferred estimate of the slope in the backwards-looking Phillips Curve of .14. Although there is uncertainty surrounding the value of slope of alternative Phillips curve specifications (for example, Gali and Gertler (1999) suggest an estimate for the backwards-looking PC slope of 0.081 while they propose a range for the NKPC in 0.015 and 0.037), we believe that fixing the slope in each Phillips Curve equation at .1 is appropriate in order to allow us to focus on the uncertainty with respect to the presence of the backwards-looking versus forward-looking elements of the state dynamics.

[22]Given that we consider the case $b_B=b_F=.1$ the choice of $\mu$ is irrelevant for the results reported. Experiments with different choices of $b_B$ and $b_F$ (and in turn of $\mu$) did not yield qualitatively different results with respect to the main messages we wish to communicate from our exercises.

Tables 1–4 present results for the optimal backwards-looking rule $O_B$, equation (14), the optimal forward-looking rule $O_F$, equation (16), the restricted optimal rule $RO_F$, equation (18), and the rule of thumb, $RoT$. The spectral densities induced by these rules are presented in Figures 1 and 2. The various loss calculations are predicated on the loss function (12) and either on the backwards-looking model (13) or the forward-looking model (15). Policymaker losses are reported for the backwards- and forward-looking models across the four policy rules in Tables 1 and 2. Tables 3 and 4 include minimax and minimax regret calculations. Besides the total variance of $\pi_t$, Tables 1-4 contain the variance of $\pi_t$ computed on several frequency intervals. The intervals are $\left[0,\frac{\pi}{4}\right]$, $\left[\frac{\pi}{4},\frac{\pi}{2}\right]$, $\left[\frac{\pi}{2},\frac{3\pi}{4}\right]$, and $\left[\frac{3\pi}{4},\pi\right]$, along with the cumulated intervals $\left[\frac{\pi}{4},\pi\right]$ and $\left[\frac{\pi}{2},\pi\right]$. These intervals reveal how frequency-specific preferences matter for policy evaluation. For example, $\left[0,\frac{\pi}{4}\right]$ gives the loss when a policymaker only cares about lower frequency fluctuations, while $\left[\frac{3\pi}{4},\pi\right]$ represents loss when a policymaker only cares about higher frequency fluctuations, and $\left[\frac{\pi}{4},\pi\right]$ and $\left[\frac{\pi}{2},\pi\right]$ indicate loss for all but the lower frequency fluctuations. These approximate the frequency intervals considered by Otrok (2001). Note that frequencies 0, $\frac{\pi}{4}$, $\frac{\pi}{2}$, $\frac{3\pi}{4}$ and $\pi$ correspond to $\infty$, 8, 4, 2.67, and 2 periods per cycle, respectively.

Consider first the performance of different rules when the backwards-looking model (13) is true. Table 1 shows, not surprisingly, that $O_B$ dominates the $O_F$, $RO_F$, and $RoT$ policy rules. However, the $RO_F$ and $RoT$ policy rules yield total variances close to $O_B$, which has a total variance equal to the normalized variance of the shock $\varepsilon_t$. In contrast, $O_F$ performs poorly generating a total variance nearly four times larger.

The frequency intervals contain information that explains the poor performance of the $O_F$ rule for the backwards-looking model (13). Across the frequency intervals

$\left[0,\frac{\pi}{4}\right]$, $\left[\frac{\pi}{4},\frac{\pi}{2}\right]$, $\left[\frac{\pi}{2},\frac{3\pi}{4}\right]$, and $\left[\frac{3\pi}{4},\pi\right]$, this rule produces variances of 0.15, 3.15, 0.28, and 0.14, which appear in the second line of Table 1. Thus, the $O_F$ rule creates large fluctuations in $\pi_t$ in the middle to low frequencies. This spike in the spectrum of $\pi_t$ appears in Figure 1, which plots spectral densities of $\pi_t$ across the four policy rules for the backwards-looking model (13). Note that the $O_B$, $RO_F$, and $RoT$ rules generate nearly flat spectra that reflect the ability of these rules to annihilate fluctuations in $\pi_t$ across all frequencies. Since the $O_F$ rule cannot eliminate movements in $\pi_t$ at the business cycle frequencies (i.e., the middle to low frequencies), which maybe the ones the policymaker cares about most, it appears costly for the policymaker's staff to rely on the forward-looking model (15) when the underlying model is not.

The reduced form of $\pi_t$ explains the behavior of $O_F$ in the backwards-looking model. Under this policy rule, $\pi_t$ is an ARMA(2,1) process with associated MA$\left(\infty\right)$ of the form

$$\pi_t = \frac{\left(1-\frac{\beta}{\delta}L\right)}{1-\left(\frac{\beta}{\delta}+a-\frac{b_B}{b_F}\frac{\beta}{\delta}\right)L+a\frac{\beta}{\delta}L^2}\varepsilon_t . \tag{19}$$

The roots of the polynomial in the denominator of (19) are

$$\frac{\left(\frac{\beta}{\delta}+a-\frac{b_B}{b_F}\frac{\beta}{\delta}\right)\pm\sqrt{\left(\frac{\beta}{\delta}+a-\frac{b_B}{b_F}\frac{\beta}{\delta}\right)^2-4a\frac{\beta}{\delta}}}{2a\frac{\beta}{\delta}} . \tag{20}$$

The discriminant $\left(\frac{\beta}{\delta}+a-\frac{b_B}{b_F}\frac{\beta}{\delta}\right)^2-4a\frac{\beta}{\delta}$ is negative given our calibrations, which yields complex conjugate roots for the polynomial in the denominator of the MA process

(19). In this case, the spectral density of the state produces a peak at business cycle frequencies.

In contrast, $RO_F$ policy produces a more parsimonious AR(1) process with associated MA representation

$$\pi_t = \frac{1}{1-\left(a+\frac{b_B}{b_F}\frac{\left(1-\sqrt{1+8\beta^2}-4\beta^2\right)}{8\beta^2}\right)L}\varepsilon_t . \tag{21}$$

No spurious cycles are induced in this case because the policymaker is able to wipe out the first-order serial correlation in $\pi_t$ at almost all frequencies. Figure 1 provides evidence that much the same holds for the $O_B$ and $RO_F$ policy rules, conditional on the loss function (12) and the backwards-looking model (13).

When one considers the frequency-specific effects of policies, some interesting differences emerge in the comparisons. The policy rules are evaluated differently across the frequency intervals, except for the mid- to low frequency interval $\left[\frac{\pi}{4},\frac{\pi}{2}\right]$. For example, the $O_F$ rule performs better than $O_B$ at the low frequencies, $\left[0,\frac{\pi}{4}\right]$, high frequencies, $\left[\frac{3\pi}{4},\pi\right]$, and the high to mid-frequency interval $\left[\frac{\pi}{2},\pi\right]$. Since time separable preferences place less weight on the low frequencies, it shows that support for the $O_B$ policy rule in the backwards-looking model (13) depends on the loss function (12). Although this result is not controversial, the superior performance of the $O_F$ policy rule in the backwards-looking model (13) at specific frequency intervals suggests that the choice of the 'best' policy relies on knowledge of the policymaker's preferences.

Table 2 and Figure 2 record results when the forward-looking model (15) is true. The $O_F$ rule dominates the $O_B$, $RO_F$, and $RoT$ policies for total variance as well as across the frequency intervals on which we have focused. Note that the $O_B$, $RO_F$, and

$RoT$ policies generate losses three to eight times larger than produced by the $O_F$ policy rule at the high frequencies. However, the superiority of $O_F$ is not as clear cut as it is for $O_B$ when the true model is the backwards-looking model (13). At the low to middle frequencies, the losses under the $O_B$, $RO_F$, and $RoT$ policies are close to those produced by the $O_F$ rule when the forward-looking model (15) is correct. Thus, model uncertainty creates an asymmetry in the conduct of policy evaluation.

We present minimax and minimax regret calculations in Tables 3 and 4 to evaluate the strength of the case for the model-specific optimal rules. The tables reveal that the preferred rule depends on the criterion used to evaluate policymaker loss. For total variance, Table 3 show that $O_B$ is favored under minimax, but minimax regret chooses the $RO_F$ rule according to Table 4.[23] This shows that when a rule performs poorly it affects policy comparisons. Given that a policymaker strives to mitigate the impact of the worst outcomes, we argue that minimax regret is preferred for policy evaluation. Unlike minimax, minimax regret is designed to account for these efforts by the policymaker.

Frequency-specific comparisons produce different conclusions. Tables 3 and 4 show that minimax never selects $O_B$, while minimax regret chooses $O_F$ most often across the frequency intervals. Under minimax, the $RoT$ policy dominates at the low to middle. At the middle to high frequencies, minimax favors either the $O_F$ or $RO_F$ rules. The $O_F$ rule is not chosen by minimax regret only at the middle to low frequency interval, where the $RoT$ policy performs best, and at the high to low frequency interval, where the $RO_F$ policy dominates. Although Tables 3 and 4 indicate that the $O_F$ rule generates the largest total variance either for minimax or for minimax regret, this rule

[23]We note that the minimax regret choice of $RO_F$ for overall variance is the one case where independence of irrelevant alternatives is violated: if $O_F$ is removed from the policy rule set, then minimax regret chooses $O_B$. The reason for this is that when $O_F$ is removed, another rule becomes the best choice among the candidate rule for the forward-looking case and so alters the regret calculations. Our comparisons all fulfill the independence of nonoptimal alternatives criterion described above.

performs best at the higher frequencies under these criteria. Moreover, minimax regret often settles on the $O_F$ rule across most of the frequency intervals. Thus, total variance can mask frequency-specific policy trade-offs.

This section shows that policy evaluation is sensitive to model and loss function uncertainty. The minimax and minimax regret exercises reveal that frequency-specific fluctuations matter most for the latter type of uncertainty. Given the backwards- and forward-looking models of equations (13) and (15), the policymaker chooses the $O_F$ rule most often when he places more weight on high frequency fluctuations under minimax or minimax regret. For total variance, the minimax and minimax regret are split on which policy rule to follow. This suggests that the benefits of optimal rules are conditional on the loss function imposed on the policymaker.

## *b.* hybrid models

This section studies a hybrid model that has backwards- and forward-looking elements. We meld equations (13) and (15) into a continuum of potential hybrid models by varying the weights on the backwards- and forward-looking components. The result is

$$\pi_t = (1-\theta)\beta E_t \pi_{t+1} + \theta a \pi_{t-1} + b(\theta) y_{t-1} + \varepsilon_t; \qquad \varepsilon_t \sim WN \tag{22}$$

where $\theta \in [0,1]$ and $b(\theta) = \theta b_B + (1-\theta) b_F$. Under the objective function (12) it can be shown that the fully optimal policy for this model under the objective function is

$$y_t^*(\theta) = \frac{\alpha(\theta)}{b(\theta)} \frac{(1-B(\theta)L)}{(1-\gamma(\theta)L)} \pi_t \tag{23}$$

where $\alpha(\theta) = -\frac{(1-\theta)\beta + \theta a\delta}{\delta}$, $B(\theta) = \frac{\theta(1-\theta)a\beta}{(1-\theta)\beta + \theta a\delta}$, and $\gamma(\theta) = (1-\theta)\frac{\beta}{\delta}$. Notice that as $\theta \to 0$ this policy converges to the policy that is optimal under the purely forward-looking model. Given this rule, a rational expectations equilibrium exists and is unique when

$$\frac{(1-\theta)\beta^2+\delta}{\beta\delta}>1 \tag{24}$$

which we assume. The process for the state variable under the optimal policy is

$$\pi_t=\frac{\delta}{(1-\theta)\beta^2+\delta}\left(1-(1-\theta)\frac{\beta}{\delta}L\right)\varepsilon_t . \tag{25}$$

Suppose now that the benchmark model is one with $\theta=\hat{\theta}$. We are interested in evaluating how this policy performs across the model space $\theta\in[0,1]$ compared to a generic rule of thumb $y_t=\bar{f}\pi_t$.

The respective laws of motion for the state variable under the rule of thumb and the model-specific optimal policies are

$$\pi_t=(1-\theta)\beta E_t\pi_{t+1}+\theta a\pi_{t-1}+b(\theta)\,\bar{f}\,\pi_{t-1}+\varepsilon_t \tag{26}$$

and

$$\pi_t=(1-\theta)\beta E_t\pi_{t+1}+\theta a\pi_{t-1}+b(\theta)\,y^*_{t-1}\left(\hat{\theta}\right)+\varepsilon_t . \tag{27}$$

respectively. For the rule of thumb, the solution for the state is an AR(1) process

$$\pi_t=\frac{1}{\lambda_1\left(\theta,\bar{f}\right)\beta(1-\theta)}\frac{1}{\left(1-\lambda_2\left(\theta,\bar{f}\right)L\right)}\varepsilon_t \tag{28}$$

where

$$\frac{1}{\lambda_i\left(\theta,\overline{f}\right)}=\frac{1\pm\sqrt{1-4(1-\theta)\beta\left(\theta a+\overline{f}b(\theta)\right)}}{2\left(\theta a+\overline{f}b(\theta)\right)}. \tag{29}$$

Existence and uniqueness require that

$$\left|\lambda_1\left(\theta,\overline{f}\right)\right|>1,\left|\lambda_2\left(\theta,\overline{f}\right)\right|<1. \tag{30}$$

These conditions hold for the calibration we study.

The model-specific optimal rule induces a law of motion for the state of the form

$$\pi_t=(1-\theta)\beta\left(1-\gamma\left(\hat{\theta}\right)L\right)E_t\pi_{t+1}+H\left(\theta,\hat{\theta}\right)\pi_{t-1}+K\left(\theta,\hat{\theta}\right)\pi_{t-2}+\left(1-\gamma\left(\hat{\theta}\right)L\right)\varepsilon_t. \tag{31}$$

in our generic hybrid model, where $H\left(\theta,\hat{\theta}\right)\equiv\gamma\left(\hat{\theta}\right)+\theta a+b(\theta)\frac{\alpha\left(\hat{\theta}\right)}{b\left(\hat{\theta}\right)}$ and $K\left(\theta,\hat{\theta}\right)=-\theta a\gamma\left(\hat{\theta}\right)-b(\theta)\frac{\alpha\left(\hat{\theta}\right)B\left(\hat{\theta}\right)}{b\left(\hat{\theta}\right)}$. These two terms are the equilibrium coefficients on past inflation resulting from the implementation and are critical in understanding how model uncertainty affects the evaluation of a model-specific optimal rule. If the policymaker works with the correct model, the backwards-looking terms in (31) can be shown to vanish, i.e. $H(\theta,\theta)=0$ and $K(\theta,\theta)=0$. When the assumed model is not correct the, the functions $H\left(\theta,\hat{\theta}\right)$ and $K\left(\theta,\hat{\theta}\right)$ are generally different from zero. The intuition for this is that optimal policy exploits the presence of forward-looking agents in the model. A commitment to the rule (23) modifies the expectations of agents by forcing them to internalize the anticipated dynamics of $\pi_t$ as modified by the policy rule. When the rule is based on the correct model, this eliminates the direct role of lags in $\pi_t$ so that changes in expectations exclusively drive the $\pi_t$ dynamics. When the policy rule is optimized based on a misspecified model, $\pi_t$ depends on its own lags, which affects its time series properties.

The moving average representation of the state variable, $\pi_t = G(L)\varepsilon_t$ can be shown to follow

$$G(L) = \frac{\left(1-\gamma\left(\hat{\theta}\right)L\right)\left(G_0\left(1-\theta\right)\beta - L\right)}{\left(\left(1-\theta\right)\beta - \left(1+\left(1-\theta\right)\beta\gamma\left(\hat{\theta}\right)\right)L + H\left(\theta,\hat{\theta}\right)L^2 + K\left(\theta,\hat{\theta}\right)L^3\right)} \tag{32}$$

when the policymaker employs a misspecified model. This representation is neither unique nor bounded unless restrictions are imposed in the roots of the autoregressive polynomial. This polynomial can be rewritten as

$$\left(1-\theta\right)\beta\left(1-\lambda_1\left(\theta,\hat{\theta}\right)L\right)\left(1-\lambda_2\left(\theta,\hat{\theta}\right)L\right)\left(1-\lambda_3\left(\theta,\hat{\theta}\right)L\right). \tag{33}$$

A sufficient condition for a unique and bounded solution is

$$\left|\lambda_1\left(\theta,\hat{\theta}\right)\right| > 1, \left|\lambda_2\left(\theta,\hat{\theta}\right)\right| < 1, \left|\lambda_3\left(\theta,\hat{\theta}\right)\right| < 1 \tag{34}$$

which we assume.[24] The solution is again an ARMA(2,1)

$$\pi_t = \frac{\left(1-\gamma\left(\hat{\theta}\right)L\right)}{\left(\lambda_1\left(\theta,\hat{\theta}\right)\left(1-\theta\right)\beta\right)\left(1-\lambda_2\left(\theta,\hat{\theta}\right)L\right)\left(1-\lambda_3\left(\theta,\hat{\theta}\right)L\right)}\varepsilon_t . \tag{35}$$

[24]The presence of $G_0$ on the right hand side of (32) is the source of potential non-uniqueness, while the possibility of a root at the denominator polynomial bigger than one in absolute value is the source of potential unboundedness. Following Whiteman (1983), the term $G_0$ is chosen so that the unstable autoregressive root $\left|\lambda_1\left(\theta,\hat{\theta}\right)\right| > 1$ is cancelled with a root of the polynomial at the numerator in (32). Once the cancellation is performed, the unique bounded solution is an ARMA (2,1) as reported in the text.

under the misspecified model. Since the autoregressive roots $\lambda_2\left(\theta,\hat{\theta}\right)$ and $\lambda_3\left(\theta,\hat{\theta}\right)$ are zero under the correct model ($\theta=\hat{\theta}$), (35) reveals how conditioning an optimal policy on the wrong model will generate fluctuations in $\pi_t$.

We conduct several numerical exercises to gauge the impact of model uncertainty on the policymaker's loss when the economic model is eq. (22) . The calibration follows the previous exercise, which sets $\beta=0.9$, $\delta=0.99$, $a=0.9$, and $b_B=b_F=.1$. The optimal rule is constructed on the basis of the benchmark model, where $\theta=0.5$, which yields the optimal rule $y_t=-0.90\dfrac{(1-0.22L)}{(1-0.45L)}\pi_t$. We continue to engage the rule of thumb $y_t=-\dfrac{1.25}{b}\pi_{t-1}$. Note that the rational expectations equilibrium exists, is unique and bounded in all the exercises discussed below.

Table 5 presents losses over the model space $\theta\in[0,1]$ for the optimal benchmark rule (*OB*) and the rule of thumb ($RoT$). The Table describes average, minimum, maximum, maximum regret losses for the rules for total variance and across the frequency intervals. Average loss is taken across the model space under the assumption that $\theta$ is uniformly distributed on the interval $[0,1]$.

Table 5 reveals how the comparisons between the model-specific optimal rule and rule of thumb depend on which frequencies are considered. Starting with average loss, the *OB* rule dominates the $RoT$ rule at the mid- to high frequencies, i.e. $\left[\frac{\pi}{2},\pi\right]$. This translates into the *OB* rule generating lower total variance. However, $RoT$ is selected over *OB* on the intervals $\left[0,\frac{\pi}{4}\right],\left[\frac{\pi}{4},\frac{\pi}{2}\right]$ and $\left[\frac{\pi}{2},\frac{3\pi}{4}\right]$. The average loss comparisons are mirrored in all cases by the minimax and minimax regret comparisons of the rules. Thus, we find that the ranking of rules is sensitive to the frequency intervals used to evaluate policy, which is consistent with results reported in the previous section.

Figure 3 transforms the total and frequency interval variances of $\pi_t$ under the *OB* and $RoT$ rules found in Table 5 into plots that move pointwise through the model space, $\theta \in [0,1]$. The dashed lines denote variance generated by the *OB* rule, while the $RoT$ variances are solid lines in Figure 3. Figure 3-a shows that the *OB* rule always yields smaller total variance pointwise across the model space. As Table 5 and the associated panels in Figure 3 suggest, the same is true on the $\left[\frac{3\pi}{4}, \pi\right]$, $\left[\frac{\pi}{2}, \pi\right]$, and $\left[\frac{\pi}{4}, \pi\right]$ frequency intervals. Also consistent with Table 5 is that on the low to mid-frequency intervals the $RoT$ almost always generates a smaller variances across the model space. Nonetheless, the *OB* rule comes close to the $RoT$ for $\theta \in [0,.4]$ on the $\left[0, \frac{\pi}{4}\right]$ and $\left[\frac{\pi}{4}, \frac{\pi}{2}\right]$ frequency intervals and on the $\left[\frac{\pi}{2}, \frac{3\pi}{4}\right]$ frequency interval at $\theta \in [.9,1]$. Contrast this with results reported in Table 1 and Figure 1 that the $RoT$ dominates the optimal forward-looking rule, $O_F$, when $\theta = 1$. The source of the disparate results is that the optimal benchmark rule *OB* is constructed under the assumption that $\theta = 0.5$. This assumption minimizes deviations from the true model and hence model uncertainty, unlike the exercises reported in Table 1 and Figure 1 that do not combine the competing models.

Model uncertainty continues to matter once policy is evaluated across the entire spectrum for a fixed $\theta$. Figure 4 displays spectral densities for $\pi_t$ that are produced by the *OB* rule and the $RoT$ at $\theta = 0.5$, 0.25, and 0.75. Although $\theta = 0.5$ minimizes variance for the *OB* rule, the top panel of Figure 4 reveals that the $RoT$ produces less power in the spectral density of $\pi_t$ on the frequency interval $\left[\frac{\pi}{4}, \frac{3\pi}{4}\right]$. At $\theta = 0.25$ and 0.75, spectral densities of $\pi_t$ are qualitatively similar across the *OB* rule and the $RoT$. In this case, the superiority of the *OB* rule is derived from its ability to reduce fluctuations in $\pi_t$ at the high frequencies. This indicates that outcomes under the *OB* rule and the $RoT$ are sensitive to the frequency intervals on which these rules are compared.

Thus, policy evaluation becomes more difficult once it is recognized that model and preference uncertainty can interact.

We conclude this section by noting that there is evidence for the policymaker to choose the $OB$ rule over the $RoT$. However, this result rests on certainty about the policymaker's loss function (i.e., total variance). Our exercises acknowledge model uncertainty along the lines of backwards- versus forward-looking elements, but this uncertainty plays a secondary role because of the simplicity of our economic model. Nonetheless, our results indicate the need for policy evaluation to account for preference uncertainty with the same effort that is dedicated to model uncertainty.

## 5. Summary and conclusions

In this paper, we have attempted to illustrate how the case for optimal policy rules is weakened when one considers the performance of the rules in the presence of model uncertainty and with respect to frequency-specific effects. These two dimensions interact as we observe that the potential for optimal rules to go awry seems especially great for particular frequency-specific intervals. This suggests to us that great caution should be taken in seriously advocating a policy that is optimal with respect to an environment, using the standard variance-based loss function. While we are not so nihilistic as to believe such knowledge is impossible, we believe that it is often the case that contemporaneous scholarly discussions of macroeconomic policy pay too little attention to Friedman's 1948 arguments. Indeed, our analysis is more interventionist than his in that our simple rules are versions of leaning against the wind policies, which Friedman specifically questions because of the problem of long and variable lags in policy effects. We do not think it unfair to say that modern time series analysis has led to a more optimistic view of the information available to policymakers than assumed by

Friedman.[25] Hence, we are comfortable with recommendations that are more interventionist than his.[26]

A major weakness in our analysis is that we focus on comparisons of permanent rule choices. In other words, we ask whether a policymaker should permanently prefer one rule or the other, without allowing for the possibility that the policymaker will learn about the nature of model uncertainty over time. While this is the standard procedure in most of the monetary rules literature, the analysis fails to address how a rule should be chosen when the policymaker has the option of changing it in response to new information that affects the model space under consideration. One solution is suggested by Svensson and Williams (2005) who treat different models as regimes across which the economy switches; in their analysis the model uncertainty facing a policymaker is equivalent to uncertainty about the regime in which the economy is currently in as the conditional probabilities of future regimes given the current regime state. This is a promising research direction and admits progress using Markov jump process methods, but it treats model uncertainty in a very different way than we conceptualize it. In our view, model uncertainty represents something that may be resolved over time, not something that exists in a steady state. Hence the next step in the approach we take seems likely to draw from ideas from the theory of bandit problems rather than from the theory of Markov jump processes. Further, both approaches to model uncertainty cannot address the issue of policy evaluation when new elements emerge in the model space over time. This problem begins to link monetary policy evaluation with issues that lie at the frontiers of the work on decisionmaking under ambiguity, where even the most advanced treatment typically assumes that while the probabilities are not available for the object of interest, its support is known.

Finally, we observe that our frequency-specific results suggest that a systematic examination of preference uncertainty is an important complement to the analysis of model uncertainty. Our findings that rules that are optimal with respect to overall variance mask frequency tradeoffs, even when there is no model uncertainty, suggests a

[25] In fact, Friedman (1972) acknowledges some progress of this type, although worry about politicization preserves his skepticism of countercyclical monetary policy.

[26] As such, we are sympathetic to the common sense spirit of Blinder (1998) though in our parlance a rule can react to the state of the economy.

new dimension along which to consider the question of robustness. Brock, Durlauf and Rondina (2007) show, for example, that in forward looking systems, one can design feedback rules that induce frequency by frequency variance reductions. The consideration of this type of property as a distinct form of robustness seems a promising direction.

Therefore, we see this paper as only a first step in a long research program.

## Table 1

### Losses for Backwards Model Under Four Alternative Policies

| | Total Variance | Low Freq. $[0, \pi/4]$ | Middle-Low $[\pi/4, \pi/2]$ | Middle-High $[\pi/2, 3\pi/4]$ | High $[3\pi/4, \pi]$ | High to Mid $[\pi/2, \pi]$ | High to Low $[\pi/4, \pi]$ |
|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
| $O_B$ | 1.00 | 0.25 | 0.25 | 0.25 | 0.25 | 0.50 | 0.75 |
| $O_F$ | 3.71 | 0.15 | 3.15 | 0.28 | 0.14 | 0.41 | 3.57 |
| $RO_F$ | 1.05 | 0.38 | 0.28 | 0.21 | 0.18 | 0.39 | 0.67 |
| $RoT$ | 1.14 | 0.14 | 0.18 | 0.30 | 0.52 | 0.81 | 1.00 |

Table 1 reports the loss measured by the unconditional variance of $\pi_t$ for the backwards-looking model of equation (13) under four alternative policy rules: $O_B$ (Optimal Backwards) – the optimal rule of the backwards-looking model (equation (14)); $O_F$ (Optimal Forward) – the optimal rule of the forward-looking model (equation (16); $RO_F$ (Restricted Optimal) – the optimal rule of the forward-looking model in the class of simple first-order feedback rules (equation (20)); the $RoT$ (Rule of Thumb) of the form $y_t = -\frac{1.25}{b}\pi_t$. Column (1) reports the loss across the entire frequency range, columns (2)-(7) report the loss generated at relevant frequency intervals. The loss generated at the interval $\omega \in [\omega_L, \omega_H]$ is $2\int_{\omega_L}^{\omega_H} f_\pi(\omega) d\omega$ where $f_\pi(\omega) = \frac{1}{2\pi}\left|G\left(e^{-i\omega}\right)\right|^2$ and $\pi_t = G(L)\varepsilon_t$ ($\varepsilon_t$ white noise, $\sigma_\varepsilon^2 = 1$).

## Table 2

### Loss for Forward Model Under Four Alternative Policies

| | Total Variance | Low Freq. $[0, \pi/4]$ | Middle-Low $[\pi/4, \pi/2]$ | Middle-High $[\pi/2, 3\pi/4]$ | High $[3\pi/4, \pi]$ | High to Mid $[\pi/2, \pi]$ | High to Low $[\pi/4, \pi]$ |
|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
| $O_B$ | 0.65 | 0.04 | 0.06 | 0.13 | 0.42 | 0.55 | 0.61 |
| $O_F$ | 0.17 | 0.00 | 0.03 | 0.06 | 0.08 | 0.14 | 0.17 |
| $RO_F$ | 0.63 | 0.06 | 0.08 | 0.15 | 0.35 | 0.50 | 0.58 |
| $RoT$ | 0.81 | 0.03 | 0.04 | 0.10 | 0.64 | 0.74 | 0.78 |

Table 2 reports the loss measured by the unconditional variance of $\pi_t$ for forward-looking model of equation (15) under four alternative policy rules: $O_B$ (Optimal Backwards) – the optimal rule of the backwards-looking model (equation (14)); $O_F$ (Optimal Forward) – the optimal rule of the forward-looking model (equation (16)); $RO_F$ (Restricted Optimal) – the optimal rule of the forward-looking model in the class of simple first-order feedback rules (equation (18)); the $RoT$ (Rule of Thumb) of the form $y_t = -\frac{1.25}{b}\pi_t$. Column (1) reports the loss across the entire frequency range, columns (2)-(7) report the loss generated at relevant frequency intervals. The loss generated at the interval $\omega \in [\omega_L, \omega_H]$ is $2\int_{\omega_L}^{\omega_H} f_\pi(\omega)d\omega$ where $f_\pi(\omega) = \frac{1}{2\pi}\left|G\left(e^{-i\omega}\right)\right|^2$ and $\pi_t = G(L)\varepsilon_t$ ($\varepsilon_t$ white noise, $\sigma_\varepsilon^2 = 1$).

## Table 3

### Minimax Calculations

| | | Total Variance | Low Freq. $[0,\pi/4]$ | Middle-Low $[\pi/4,\pi/2]$ | Middle-High $[\pi/2,3\pi/4]$ | High $[3\pi/4,\pi]$ | High to Middle $[\pi/2,\pi]$ | High to Low $[\pi/4,\pi]$ |
|---|---|---|---|---|---|---|---|---|
| | | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
| Maximum | $O_B$ | **1.00** | 0.25 | 0.25 | 0.25 | 0.42 | 0.55 | 0.75 |
| | $O_F$ | 3.71 | 0.15 | 3.15 | 0.28 | **0.14** | **0.41** | 3.57 |
| | $RO_F$ | 1.05 | 0.38 | 0.28 | **0.21** | 0.35 | 0.50 | **0.67** |
| | $RoT$ | 1.14 | **0.14** | **0.18** | 0.30 | 0.64 | 0.81 | 1.00 |
| Minimax | | $O_B$ | $RoT$ | $RoT$ | $RO_F$ | $O_F$ | $O_F$ | $RO_F$ |

Table 3 reports the maximum losses across the backwards- and forward-model models from the entries reported in Table 1 and in Table 2 for the four alternative policy rules $O_B$, $O_F$, $RO_F$ and $RoT$. Column (1) reports the maximum loss across the entire frequency range, columns (2)-(7) report the maximum loss generated at relevant frequency intervals. The entries in bold indicate the minima of the maximum losses in each of the columns. The bottom row of Table 3 reports the minimax policy strategies, i.e. the rules that would ensure at most the minimum of the maximum losses.

## Table 4

### Minimax Regret Calculations

| | | Total Variance | Low Freq. $[0, \pi/4]$ | Middle-Low $[\pi/4, \pi/2]$ | Middle-High $[\pi/2, 3\pi/4]$ | High $[3\pi/4, \pi]$ | High to Middle $[\pi/2, \pi]$ | High to Low $[\pi/4, \pi]$ |
|---|---|---|---|---|---|---|---|---|
| | | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
| | | | | **Backward** | | | | |
| Regret | $O_B$ | 0.00 | 0.11 | 0.07 | 0.04 | 0.11 | 0.11 | 0.08 |
| | $O_F$ | 2.71 | 0.00 | 2.97 | 0.07 | 0.00 | 0.03 | 2.90 |
| | $RO_F$ | 0.05 | 0.23 | 0.10 | 0.00 | 0.04 | 0.00 | 0.00 |
| | $RoT$ | 0.14 | 0.00 | 0.00 | 0.09 | 0.38 | 0.43 | 0.33 |
| | | | | **Forward** | | | | |
| Regret | $O_B$ | 0.48 | 0.04 | 0.03 | 0.07 | 0.34 | 0.41 | 0.44 |
| | $O_F$ | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 |
| | $RO_F$ | 0.46 | 0.05 | 0.05 | 0.09 | 0.27 | 0.36 | 0.41 |
| | $RoT$ | 0.64 | 0.03 | 0.02 | 0.04 | 0.55 | 0.59 | 0.61 |
| Max Regret | $O_B$ | 0.48 | 0.11 | 0.07 | 0.07 | 0.34 | 0.41 | 0.44 |
| | $O_F$ | 2.71 | **0.00** | 2.97 | **0.07** | **0.00** | **0.03** | 2.90 |
| | $RO_F$ | **0.46** | 0.23 | 0.10 | 0.09 | 0.27 | 0.36 | **0.41** |
| | $RoT$ | 0.64 | 0.03 | **0.02** | 0.09 | 0.55 | 0.59 | 0.61 |
| Minimax Regret | | $RO_F$ | $O_F$ | $RoT$ | $O_F$ | $O_F$ | $O_F$ | $RO_F$ |

Table 4 reports losses for the regret defined by equation (6) using the backwards- and forward-looking models under four alternative policy rules: $O_B$ (Optimal Backwards) – the optimal rule of the backwards-looking model (equation (14)); $O_F$ (Optimal Forward) – the optimal rule of the forward-looking model (equation (16); $RO_F$ (Restricted Optimal) – the optimal rule of the forward-looking model in the class of simple first-order feedback rules (equation (18)); the $RoT$ (Rule of Thumb) of the form $y_t = -\frac{1.25}{b}\pi_t$. The regrets are computed from the entries in Table 1 and Table 2. Column (1) reports the maximum regret from the losses across the entire frequency range, columns (2)-(7) report the maximum regret from the losses generated at relevant frequency intervals. The entries in bold indicate the minima of the maximum regrets in each of the columns. The bottom part of Table 4 reports the minimax-regret policy strategies, i.e. the rules that would ensure at most the minimum of the maximum regrets.

## Table 5

### Losses Over the Model Space $\theta \in [0,1]$

| | | Total Variance | Low Freq. $[0, \pi/4]$ | Middle -Low $[\pi/4, \pi/2]$ | Middle -High $[\pi/2, 3\pi/4]$ | High $[3\pi/4, \pi]$ | High to Middle $[\pi/2, \pi]$ | High to Low $[\pi/4, \pi]$ |
|---|---|---|---|---|---|---|---|---|
| | | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
| | Average Loss | 0.84 | 0.08 | 0.18 | 0.25 | 0.33 | 0.58 | 0.75 |
| $OB$ | Min Loss | 0.61 | 0.03 | 0.06 | 0.15 | 0.21 | 0.50 | 0.57 |
| | Max Loss | **1.05** | 0.18 | 0.37 | 0.32 | **0.37** | **0.62** | **0.88** |
| | Max Regret | **0.00** | 0.04 | 0.19 | 0.09 | **0.00** | **0.00** | **0.00** |
| | Average Loss | 0.99 | 0.07 | 0.09 | 0.19 | 0.83 | 0.83 | 0.92 |
| $RoT$ | Min Loss | 0.81 | 0.03 | 0.04 | 0.10 | 0.74 | 0.74 | 0.78 |
| | Max Loss | 1.14 | **0.14** | **0.18** | **0.30** | 0.88 | 0.88 | 1.01 |
| | Max Regret | 0.21 | **0.00** | **0.00** | **0.01** | 0.34 | 0.32 | 0.21 |
| Minimax | | $OB$ | $RoT$ | $RoT$ | $RoT$ | $OB$ | $OB$ | $OB$ |
| Minimax Regret | | $OB$ | $RoT$ | $RoT$ | $RoT$ | $OB$ | $OB$ | $OB$ |

Table 5 reports the loss measured by the unconditional variance of $\pi_t$ for the hybrid model of equation (22) under two alternative policy rules: $OB$ (Optimal Benchmark) – the optimal rule of the benchmark model at $\hat{\theta} = 0.5$; the $RoT$ (Rule of Thumb) specification $y_t = -\frac{1.25}{b}\pi_t$. The top 8 lines report the average loss, the minimum loss, the maximum loss and the maximum regret $\max_{\theta \in [0,1]} R(p, d, \theta)$, for $R$ as defined in (6) across the model space $\theta \in [0,1]$. The bold entries denote the minima of the maximum losses and of the maximum regrets across the two policies. The last two rows report the minimax policy strategy and the minimax regret policy strategy, respectively. Column (1) reports the loss and maximum regret measured with respect to the entire frequency range, columns (2)-(7) report the loss and maximum regret measured with respect to relevant frequency intervals. The loss generated at the interval $\omega \in [\omega_L, \omega_H]$ is $2\int_{\omega_L}^{\omega_H} f_\pi(\omega) d\omega$ where $f_\pi(\omega) = \frac{1}{2\pi}\left|G\left(e^{-i\omega}\right)\right|^2$ and $\pi_t = G(L)\varepsilon_t$ ($\varepsilon_t$ white noise, $\sigma_\varepsilon^2 = 1$).

**Figure 1**

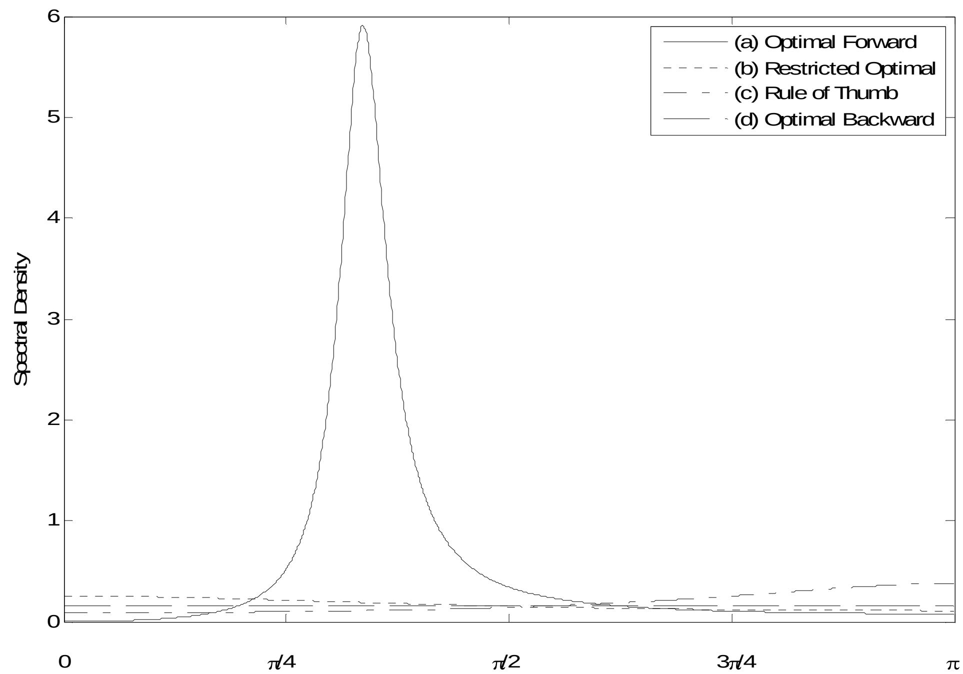


Figure 1 reports the spectral density $f_\pi(\omega) = \frac{1}{2\pi}\left|G\left(e^{-i\omega}\right)\right|^2$ for the unconditional variance of $\pi_t = G(L)\varepsilon_t$ ($\varepsilon_t$ white noise, $\sigma_\varepsilon^2 = 1$), when the backwards-looking model of equation (13) is true under four alternative policy rules: (a) $O_F$ (Optimal Forward) – the optimal rule of the forward-looking model of equation (16); (b) $RO_F$ (Restricted Optimal) – the optimal rule of the forward-looking model in the class of simple first-order feedback rules, equation (18); (c) the $RoT$ (Rule of Thumb) of the form $y_t = -\frac{1.25}{b}\pi_t$; (d) $O_B$ (Optimal Backwards) – the optimal rule under the backwards-looking model, equation (14).

**Figure 2**

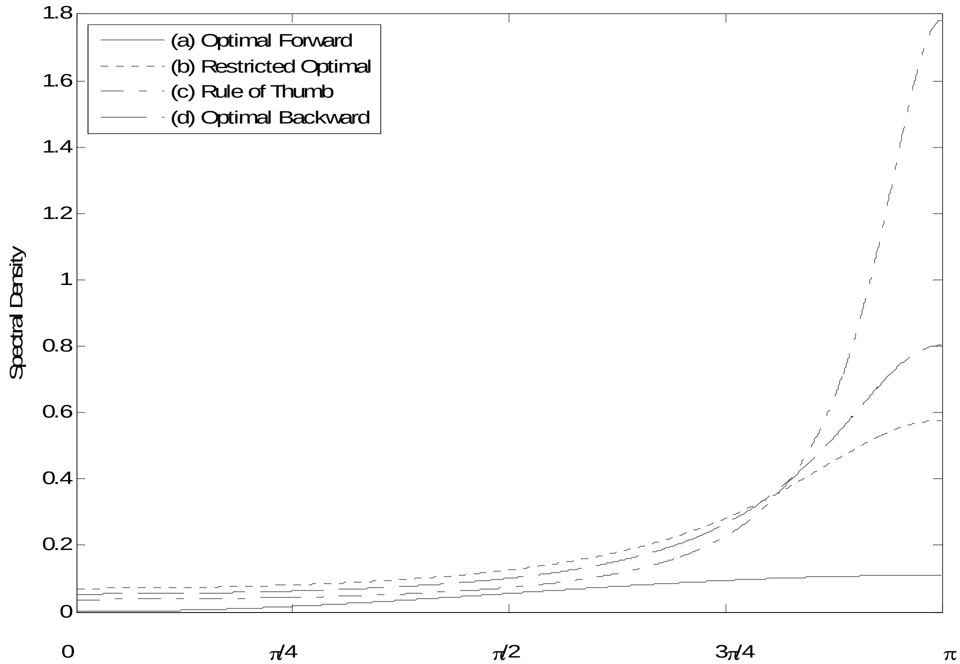


Figure 2 reports the spectral density $f_{\pi}(\omega) = \frac{1}{2\pi}\left|G\left(e^{-i\omega}\right)\right|^2$ for the unconditional variance of $\pi_t = G(L)\varepsilon_t$ ($\varepsilon_t$ white noise, $\sigma_{\varepsilon}^2 = 1$), when the model is of the forward-looking type (equation (15)) under the four alternative policy rules: (a) $O_F$ (Optimal Forward) – the optimal rule of the forward-looking model of equation (16); (b) $RO_F$ (Restricted Optimal) – the optimal rule of the forward-looking model in the class of simple first-order feedback rules, equation (18); (c) the $RoT$ (Rule of Thumb) of the form $y_t = -\frac{1.25}{b}\pi_t$; (d) $O_B$ (Optimal Backwards) – the optimal rule under the backwards-looking model, equation (14).

**Figure 3**

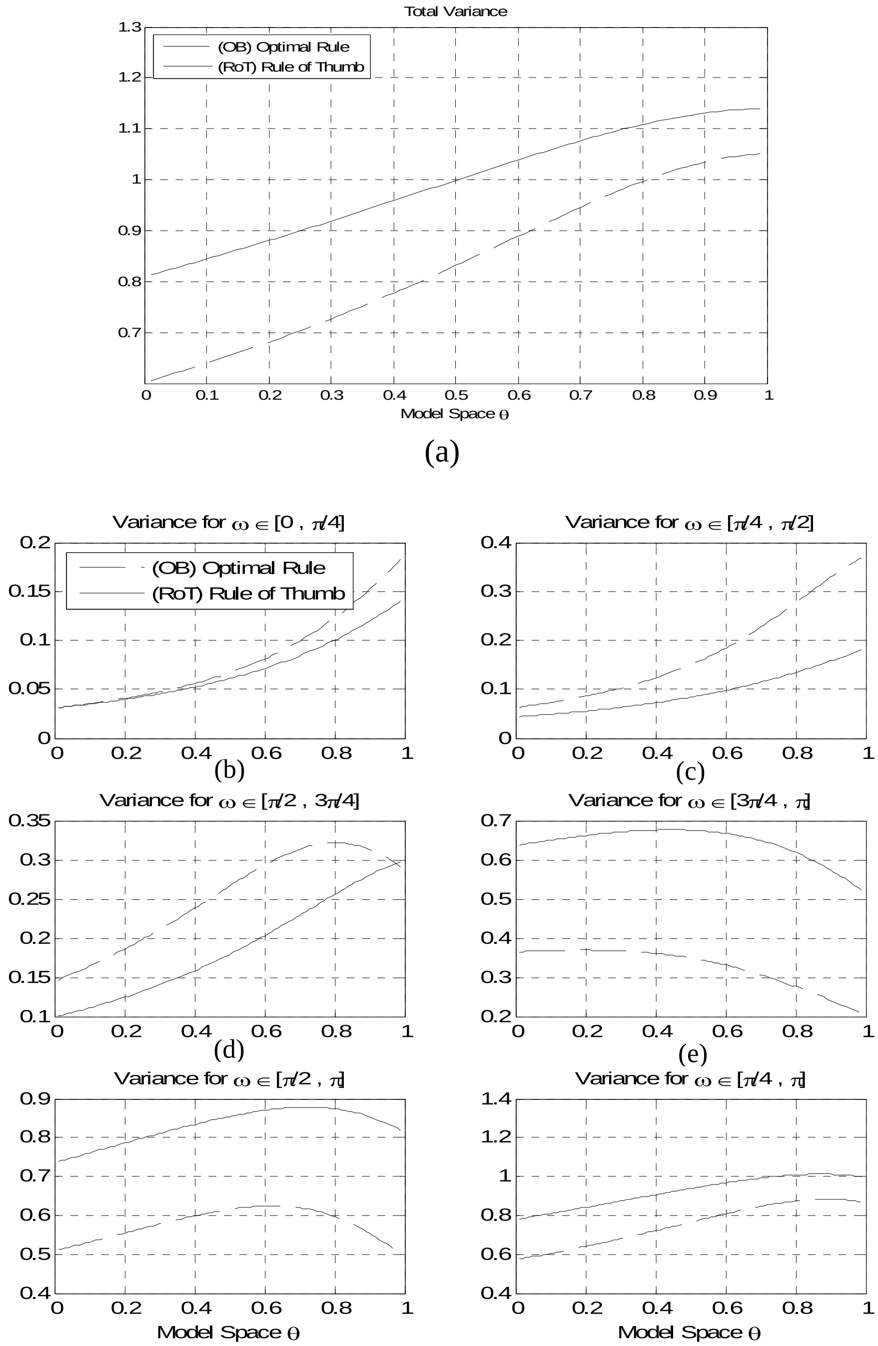


The top panel of Figure 3 reports the total unconditional variance for $\pi_t$ for the hybrid model of equation (22) across the model space, $\theta \in [0,1]$. The unconditional variance is computed under two alternative policy rule: (a) *OB* (Optimal Benchmark Rule) – the policy rule that is optimal for the benchmark model $\theta = \hat{\theta}$ where $\hat{\theta} = .5$; (b) the $RoT$ (Rule of Thumb) of the form $y_t = -\left(1.25 / b\left(\hat{\theta}\right)\right)\pi_t$. The bottom panels of Figure 3 reports the portions of the overall unconditional variance for $\pi_t$ (under the same alternative policy rules) generated at relevant frequency intervals.

# Figure 4

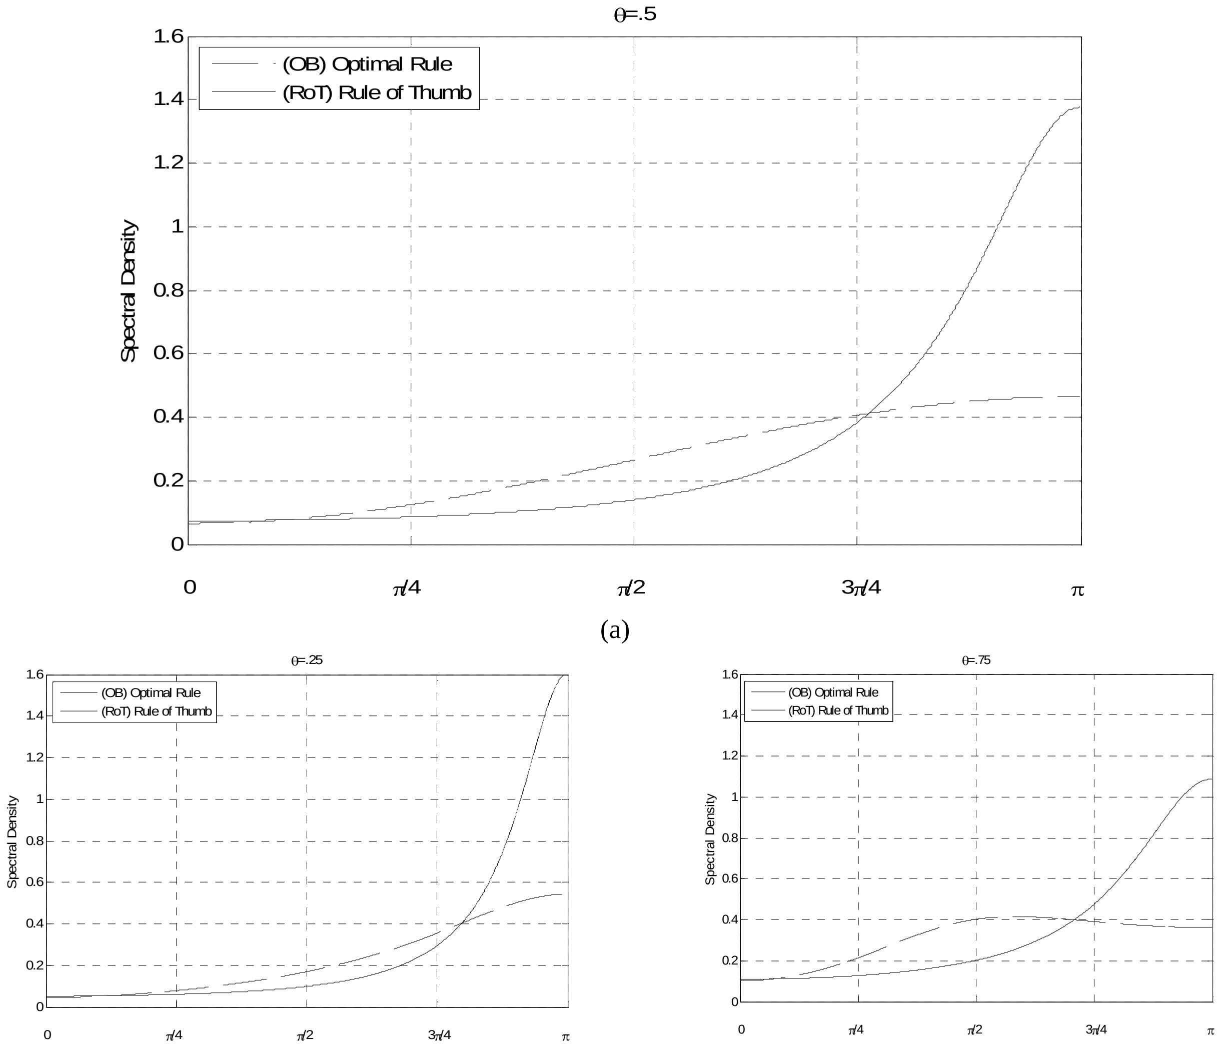


(a)

(b)

(c)

Figure 4 reports the spectral densities $f_{\pi}(\omega)=\frac{1}{2\pi}\left|G\left(e^{-i\omega}\right)\right|^{2}$ for the unconditional variance of $\pi_t = G(L)\varepsilon_t$ ($\varepsilon_t$ white noise, $\sigma_{\varepsilon}^{2}=1$), for the hybrid model of equation (22). The top panel reports the spectral densities for the model $\theta=.5$ under two alternative policy rules: (a) *OB* (Optimal Benchmark Rule) – the policy rule that is optimal for the benchmark model $\theta=\hat{\theta}$ where $\hat{\theta}=.5$; (b) the $RoT$ (Rule of Thumb) of the form $y_t=-\frac{1.25}{b}\pi_t$. The bottom left and bottom right panels report spectral densities for the two alternative policy rules given the models are $\theta=.25$ and $\theta=.75$, respectively.

**Bibliography**

Blinder, A. 1998. *Central Banking in Theory and Practice*. Cambridge: MIT Press.

Blume, L., D. Easley, and J. Halperin. 2006. Redoing the Foundations of Decision Theory. Mimeo, Cornell University.

Brock. W. 2006. Profiling Problems with Partially Identified Structure. *Economic Journal*, 116, 515, F427–F440.

Brock, W. and S. Durlauf. 2004. Elements of a Theory of Design Limits to Optimal Policy. *The Manchester Schoo*l, 72, Supplement 2, 1-18, 2004.

Brock, W. and S. Durlauf. 2005. Local Robustness Analysis: Theory and Application. *Journal of Economic Dynamics and Control* 29, 11, 2067-2092.

Brock, W., S. Durlauf, J. Nason, and G. Rondina, (2007), "Trend Uncertainty and Policy Rule Comparisons," mimeo in progress.

Brock, W., S. Durlauf, and G. Rondina. 2007. Design Limits and Dynamic Policy Analysis. Mimeo, University of Wisconsin.

Brock, W., S. Durlauf and K. West. 2003. Policy Analysis in Uncertain Economic Environments (with discussion). *Brookings Papers on Economic Activity* 1, 235-322.

Brock, W., S. Durlauf, and K. West. 2006. Model Uncertainty and Policy Evaluation: Some Theory and Empirics. *Journal of Econometrics,* forthcoming.

Chamberlain, G. 2001. Econometrics and Decision Theory. *Journal of Econometrics* 95, 255-83.

Chernoff, H. 1954. Rational Selection of Decision Functions. *Econometrica*, 22, 422-443

Clarida, R., J. Gali and M. Gertler, 1999. The Science of Monetary Policy: A New Keynesian Perspective. *Journal of Economic Literature* XXXVII(December), 1661-1701

Eozenou, P., J. Rivas, and K. Schlag, (2006), "Minimax Regret in Practice-Four Examples on Treatment Choice," mimeo, European University Institute.

Epstein, L. and T. Wang. 1994. Intertemporal Asset Pricing Behavior Under Knightian Uncertainty. *Econometrica*, 62, 283-322.

Freedman, D. 1991. A Rejoinder to Berk, Blalock, and Mason. *Sociological Methodology* 21, 353-358, 1991.

Friedman, M. 1948. A Monetary and Fiscal Framework for Economic Stability. *American Economic Review* 38, 3, 245-264.

Friedman, M. 1972. Monetary Policy. *Proceedings of the American Philosophical Society* 116, 3, 183-196.

Gali, J. and M. Gertler. (1999). Inflation Dynamics: A Structural Econometric Analysis. *Journal of Monetary Economics*, 44, 195-222.

Giannoni, M. 2002. Does Model Uncertainty Justify Caution? Robust Optimal Monetary Policy in a Forward-Looking Model. *Macroeconomic Dynamics,* 6, 111-144.

Gilboa, I. and D. Schmeidler. 1989. Maximin Expected Utility with Non-Unique Priors. *Journal of Mathematical Economics*, 18, 141-53.

Hansen, L. and T. Sargent. 2003. Robust Control of Forward-Looking Models. *Journal of Monetary Economics*, 50, 581-604.

Hansen, L. and T. Sargent. 2007. *Robustness*. manuscript, Hoover Institution, Stanford University and forthcoming, Princeton University Press.

Key, J., L. Perrichi, and A. Smith. 1998. Choosing Among Models when None of Them are True. In *Proceedings of the Workshop on Model Selection, Special Issue of Rassegna di Metodi Statistici ed Applicazioni*, W. Racugno, ed., Bologna: Pitagora Editrice.

Leamer, E. 1978. *Specification Searches*. New York: John Wiley and Sons.

Levin, A. and J. Williams. 2003. Robust Monetary Policy with Competing Reference Models. *Journal of Monetary Economics*, 50, 5, 945-975.

Manski, C. 2005. Statistical Treatment Rules for Heterogeneous Populations. *Econometrica*, 72, 1221-1246.

Manski, C. 2006a. Search Profiling With Partial Knowledge of Deterrence. *Economic Journal* 116, 515, F385–F401.

Manski, C. 2006b. Minimax-Regret Treatment Choice with Missing Outcome Data. *Journal of Econometrics*, forthcoming.

McCallum, B. 1988. Robustness Properties of a Rule for Monetary Policy. *Carnegie-Rochester Conference Series on Public Policy,* 29, 173-204.

McCallum, B. 1999. Issues in the Design of Monetary Policy Rules. *Handbook of Macroeconomics*, J. Taylor and B. McCallum, eds. Amsterdam: North-Holland, 1483-1530.

McCallum, B. and E. Nelson. 2005. Targeting versus Instrument Rules for Monetary Policy. *Federal Reserve Bank of St. Louis Review*, September/October, 597-611.

Onatski, A. and J. Stock. 2002. "Robust Monetary Policy Under Model Uncertainty in a Small Model of the U.S. Economy." *Macroeconomic Dynamics* 6: 85-110.

Onatski, A. and N. Williams. 2003. Modeling Model Uncertainty. *Journal of the European Economic Association*, 1, 1078-1122.

Otrok. C. 2001. Spectral Welfare Cost Functions. *International Economic Review*, 42, 345-367.

Raftery, A., D. Madigan, and J. Hoeting. 1997. Bayesian Model Averaging for Linear Regression Models. *Journal of the American Statistical Association*, 92, 437, 179-91.

Rudebusch, G. and L. Svensson. 1999. Policy Rules for Inflation Targeting. In *Monetary Policy Rules*, J. Taylor, ed. Chicago: University of Chicago Press.

Rudd, J. and K. Whelan. 2005. "New Tests of the new-Keynesian Phillips Curve," *Journal of Monetary Economics* 52, 1167-1181.

Savage, L. J. 1951. The Theory of Statistical Decision. *Journal of the American Statistical Association*, 46, 55-67.

Stoye, J. 2006. Statistical Decisions under Ambiguity. Mimeo, New York University.

Svensson. L. 2003. What is Wrong with Taylor Rules? Using Judgment in Monetary Policy through Targeting Rules. *Journal of Economic Literature* 42, 426-477.

Svensson, L. and N. Williams. 2005. Monetary Policy with Model Uncertainty: Distribution Forecast Targeting. Mimeo, Princeton University.

Taylor, J. 1993. Discretion Versus Policy Rules in Practice. *Carnegie-Rochester Conference Series on Public Policy* 39, 195-214.

Tetlow, R. and P. von zur Muehlen. 2001. Robust Monetary Policy with Misspecified Models: Does Model Uncertainty Always Call for Attenuated Policy? *Journal of Economic Dynamics and Control* 25, 6-7, 911-949.

Tobin, J. 1983. Monetary Policy: Rules, Targets, and Shocks. *Journal of Money, Credit, and Banking* 15, 4, 506-518.

Wald, A., 1950. *Statistical Decision Functions*. New York: John Wiley.

Whiteman, C. 1983. *Linear Rational Expectation Models: A User's Guide*. Minneapolis: University of Minnesota Press.

Woodford, M. 2003. *Interest and Prices.* Princeton: Princeton University Press.