# Local Robustness Analysis: Theory and Application

William A. Brock and Steven N. Durlauf

Revised: November 5, 2004

JEL Classification Codes: C52, E6
Keywords: robustness, model uncertainty, optimal policy, design limits

The authors thank the National Science Foundation, John D. and Catherine T. MacArthur Foundation, Vilas Trust, Institute for Research on Poverty and University of Wisconsin for financial support and Ethan Cohen-Cole and Giacomo Rondina for excellent research assistance. We have benefited from the comments of two anonymous referees.

# Local Robustness Analysis: Theory and Application

## Abstract

This paper develops a general framework for conducting local robustness analysis. By local robustness, we refer to the calculation of control solutions that are chosen so as to be optimal against the least favorable model within a small set of possible models. This set is defined local to an initial baseline model. We provide Nash and Stackelberg equilibrium characterizations of the choice of control in such contexts. We then apply this abstract formulation to the analysis of how a desire for robustness influences the choice of control for discrete time control problems of the type often found in macroeconomics. This analysis is conducted using frequency domain methods and is shown to involve certain fundamental limits to the efficacy of controls in such environments. Finally, we use these methods to identify some implications for the robust design of monetary policy.

William A. Brock
Department of Economics
University of Wisconsin
1180 Observatory Drive
Madison, WI 53706-1393
wbrock@ssc.wisc.edu

Steven N. Durlauf
Department of Economics
University of Wisconsin
1180 Observatory Drive
Madison, WI 53706-1393
sdurlauf@ssc.wisc.edu

## 1. Introduction

This paper develops elements of a theory of local robustness analysis for economic models. Since the seminal work of Hansen and Sargent (2001,2003a,2003b), there is increasing interest in the analysis of economic behavior when agents face uncertainty about the structure of the environment in which they operate. Such uncertainty has important implications for individual decisionmaking and aggregate economic outcomes as well as for policy design. Important contributions to this research program include Giannoni (2002), Marcellino and Salmon (2002), Onatski and Stock (2002), Onatski and Williams (2003) and Tetlow and von zur Muehlen (2001).

Robustness analysis differs from standard analyses of behavior under uncertainty in two respects. First, robustness analysis typically focuses on model uncertainty as opposed to uncertainty induced by the dependence of outcomes on the realizations of stochastic processes. Model uncertainty means that economic actors cannot fulfill the sorts of rationality assumptions that underlie much of modern macroeconomics, i.e. assumptions that imply that agents know the underlying structure of the economy. Second, robustness analysis, rather than transforming model uncertainty into a standard decision under uncertainty problem by placing nontrivial priors on the space of models, employs minimax-type methods to alternative actions by an agent. The minimax approaches allow one to treat robustness analysis as a zero sum noncooperative game between a policy and an opponent whom we refer to as the adversarial agent whose chooses that model within the model space that maximizes the policymaker's loss. This alters the objective of policy design away from the construction of optimal rules towards good rules, i.e. rules that perform well across all the models that are possible.

In our analysis, we will follow the local uncertainty approach that is employed in the robustness literature. We assume that the model space over which the adversarial agent maximizes against the policymaker is described as a small neighborhood around a baseline model. The radius of this neighborhood is assumed to be sufficiently small that a first-order Taylor series expansion well approximates the functions that map the radius of the model space to the decisions of the policymaker and the adversarial agent. This expansion is constructed around the decisions of the two agents that occur when the radius is zero, i.e. when there is no model uncertainty. Hansen and Sargent also work with local model uncertainty, but do so in the context of recursive structures, primarily linear quadratic ones. Hence, they are able to obtain closed form solutions (of the Ricatti) equation to the associated zero-sum game.

Relative to the existing robustness literature, we attempt to accomplish three things. First, we develop a theory of local robustness that allows one to use elementary calculus tools. Specifically, we consider model spaces in which individual elements are determined by particular values for a vector of parameters. By considering variations of this vector that are small (in the usual norm), we show that one can use basic calculus tools to derive robust analogs to optimal control solutions.[1] We are able to provide explicit characterizations of how the differences between these robust solutions and standard solutions are determined by the second derivatives of the payoff function of a policymaker. Second, we apply this general framework to the analysis of optimal feedbacks in a linear quadratic control context and derive implications of a preference for robustness in the frequency domain. This approach develops a link between robustness analysis and the existence of design limits in optimal control; see Brock and Durlauf (2004) for a discussion of design limits. We deduce some general implications of robustness for how a policymaker assesses uncertainty in the law of motion for the system and how the policymaker robustifies his policy rule in the presence of this uncertainty. In addition, we make some observations on how our framework relates to monetary policy design.

The introduction of model uncertainty considerations into macroeconomics is, in our judgment, an important development and properly reflects the continuing lack of consensus among economists on issues both of appropriate microeconomic assumptions as well as the inherent difficulties in identifying details of model structure (such as lag lengths or possible nonlinearities) once one has specified a particular theoretical framework. Robustness analysis does not represent the only way of addressing model

---

[1] See Svensson (2000) for an analysis that provides a "simple" way of understanding robustness in the context of some important macroeconomic models. Our work complements this analysis by focusing on properties of robust policies that hold for general environments. Of course, our analysis loses some of the economic substance found in Svensson's work.

uncertainty. Model uncertainty may also be addressed via standard Bayesian methods in statistical decision theory, so that the evaluation of expected policy outcomes treats model uncertainty as one of the different sorts of uncertainty that determine the distribution of outcomes given a policy. This approach is explored for macroeconomic contexts in Brock, Durlauf, and West (2003,2004) and Levin and Williams (2003). Such approaches require methods of calculating posterior model probabilities, i.e. probabilities that a given model is the "true" one conditional on prior information and available data. As such, this requires the analyst to be able to assign prior probabilities to alternative models. These sorts of calculations are explicitly avoided in the robustness approach. The assignment of priors is, of course, problematic. One appealing feature of robustness analysis is that it avoids making arbitrary probability assignments.

This is not to say that robustness analysis represents a good substitute for Bayesian approaches in all circumstances. As argued in Brock, Durlauf, and West (2003), robustness analysis is probably most appropriate, compared to the statistical decision theory approach, when the model space constitutes a set of models that are all local to some baseline model, as is done in the Hansen and Sargent work.[2] The problem with model spaces that contain very different (i.e. non-mutually local) models is that one highly improbable model can, under the minimax criterion, determine the outcome of an analysis, regardless of how improbable it is. Put differently, we believe that the importance of accounting for differences in posterior model probabilities depends on context.[3]

---

[2] Sims (2001) criticizes robust control methods for failing to address non-local model uncertainty. Our view is that the appropriateness of robust approaches to model uncertainty depends on context, so that Sims' criticisms will in some cases certainly be salient.

[3] An interesting alternative to our approach may be provided by the use of the minimax regret criterion. Manski (2004) argues that this criterion is preferable to the minimax criterion as it allows the data to implicitly affect the way that models are weighted. Specifically, Manski (2004, pg. 1228) argues that minimax can choose "treatment zero…regardless of the data." The local formulation of the minimax rule, by constraining it to a neighborhood of a baseline, allows data to influence policy choices so long as the variation of the data helps determine the neighborhood over which model uncertainty exists.

We therefore consider the local and non-local approaches to model uncertainty to be complements, not substitutes. For example, our local approach is appealing for applications where parameters estimated from a sample of size *N* describes the baseline and asymptotic theory instructs us on how fast the radius of model space declines to zero with sample size. It may also apply to contexts where one wishes to import results from a well defined baseline situation to a new situation where there is a measure of "nearness" of the new environment to the old. Since interpretations and formulations of robustness, ambiguity aversion, etc. are still controversial (e.g. see the discussion in Brock, Durlauf, and West, 2003) it seems wise to pursue several approaches at this stage of research.

This paper is organized as follows. Section 2 provides our general framework for studying local robustness. We consider Nash and Stackelberg equilibria. In addition, we discuss some aspects to robustness that arise when there are multiple discrete alternative models. Section 3 develops an analysis of local robustness for univariate discrete time systems. We develop formulas that illustrate how an adversarial agent will choose a least favorable model in such an environment and provide a parametric example. Section 4 makes some suggestions of how our analysis might inform a monetary authority. Section 5 contains summary and conclusions.

## 2. Basic theory

In this section, we describe a general framework for local robustness analysis. By local robustness analysis, we refer to the idea that a policymaker may not know the "true" model for some outcome of interest, but may have sufficient information to identify a space of potential models that is local to an initial baseline model. This may be regarded as a conservative approach to introducing model uncertainty into policy analysis, in that we start with a standard problem (identification of an optimal policy given a particular economic model) and extend the analysis to a local space of models, one that is defined by proximity to this initial baseline. The local model uncertainty assumption, in our judgment, is naturally associated with minimax approaches to policy evaluation. When a model space includes non-local alternatives, we would argue that one needs to account for posterior model probabilities in order to avoid implausible models from determining policy choice.

We consider a policymaker who wishes to minimize a loss function $J\left(u,\overline{a}\right)$ where $u \in R^n$ is a vector of control variables and $\overline{a} \in R^m$ is a vector of parameters. In our context, $J\left(u,\overline{a}\right)$ is what is referred to as a model. Standard optimal control for the policymaker may be described simply as

$$\text{minimize } J\left(u,\overline{a}\right) \text{ over } u \in R^n . \tag{1}$$

To understand robust analogs to this standard problem, we analyze this same problem under the assumption that the parameter vector is known only up to an $m-$dimensional ball around the baseline parameter value $\overline{a}$, i.e. the agent faces an unknown parameter vector $a$ where the only available prior information is that the parameter vector lies in the set defined by

$$\left\|a-\overline{a}\right\| \leq \varepsilon \tag{2}$$

where $\|\cdot\|$ denotes Euclidean distance. Each value of $a$ may thus be interpreted as indexing a different model $J\left(u,a\right)$.

Robustness analysis does not resolve uncertainty over $a$ by assigning a set of probabilities to the interval in which it lies; rather, the choice of control variables is assessed relative to the least favorable value of $a$. This approach to the choice of controls may, as argued by Hansen and Sargent, be interpreted as the search for rules that are assured to work reasonably well regardless of the actual model the agent faces out of the space of possibilities. Robustness analysis thus asks how the introduction of model uncertainty and minimax behavior alter the behavior of an agent relative to the case where the model is known.

As such, and following ideas due to Wald (1950), robustness analysis may be interpreted as the analysis of a noncooperative game in which the optimizer faces an opponent, the adversarial agent, whose objective is to maximize the optimizer's loss. This equivalence between minimax analysis and a noncooperative game between the policymaker and an adversarial agent has been exploited in the seminal work of Hansen and Sargent. The strategy space of the adversarial agent is the space of possible models and thus is represented by eq. (2). We first consider a Nash equilibrium approach to this problem; a Stackelberg equilibrium approach is considered below.

For a Nash equilibrium, the equilibrium control vector $u^*$ and parameter vector $a^*$ must fulfill two conditions. For the policymaker, it must be the case that if the adversarial agent chooses $a^*$,

$$J\left(u^*,a^*\right)\le J\left(u,a^*\right)\ \forall\ u\in R^n. \tag{3}$$

Conversely, it must be the case for the adversarial agent that if the policymaker chooses $u^*$, the adversarial agents's decision fulfills

$$J\left(u^*,a^*\right)\ge J\left(u^*,a\right)\ \forall\ a\in R^m,\ \left\|a-\overline{a}\right\|\le\varepsilon. \tag{4}$$

As indicated by (4), the adversarial agent chooses the least favorable model among those in the model space defined by (2).

In order to understand the differences in policy that are induced when one moves from optimal control of a single model to robust control over a model space, we calculate two derivatives, $\frac{du^*}{d\varepsilon}$ and $\frac{da^*}{d\varepsilon}$, when $\varepsilon=0$. This allows us to consider the effects of local model uncertainty when no uncertainty is treated as a baseline. In the subsequent analysis, we make the following assumptions:

A1. There exists a family of Nash equilibria $\{u^*\left(\varepsilon\right),a^*\left(\varepsilon\right)\}\ \varepsilon\in N$, where $N$ is a small neighborhood of 0.

A2. $J$ is $C^2$.

A3. $J_a(u,a) \neq 0$ for $(u,a)$ in an open neighborhood of $\left(u^*(0), a^*(0)\right)$.

A4. $\exists$ a bounded vector $\kappa$, which depends on $a$, such that if $|u| > \kappa$, then $J_{u_i}(u,a) > 0\ \forall i$.

A5. $J_{uu}(u,a)$ is strictly positive definite in an open neighborhood of $\left(u^*(0), a^*(0)\right)$.

A6. $\lim_{\varepsilon \to 0}\{u^*(\varepsilon), a^*(\varepsilon)\} = \{u^*(0), a^*(0)\}$.

Condition A1 means that we will avoid dealing with issues of existence of an equilibrium. Condition A2 allows us to use elementary calculus methods to analyze robustness. Condition A3 ensures that the marginal effect of a change in the adversarial agent's choice can reduce the payoff to the policymaker at equilibrium and that the adversarial agent's choice is at the boundary of his constraint set for small enough $\varepsilon$. Conditions A4 and A5 ensure that the optimal control for the policymaker is bounded and that the choice $u^*$ fulfills the first order necessary condition for optimality with equality. Condition A6 means that we assume that the Nash equilibria in our model are continuous in $\varepsilon$.

The purpose of these conditions is to impose a certain degree of regularity on the optimization problem which will facilitate calculations. In particular, what matters for our analysis is that the first-order conditions of the policymaker that determine $u^*$ are assumed to hold with equality and that the adversarial agent's choice $a^*$ lies at the boundary of his constraint set for $\varepsilon$ sufficiently near 0. None of the assumptions we employ should be regarded as especially onerous; all are made for technical convenience.

Our first result provides formulas for the optimal decisions of the policymaker and the adversarial agent under a Nash equilibrium.

**Proposition 1: Nash equilibrium robust decisions of a policymaker and an adversarial agent**

Suppose that a policymaker chooses $u$ in order to minimize $J\left(u,a\right)$ whereas an adversarial agent chooses $a$ subject to (4) in order to maximize $J\left(u,a\right)$. Under assumptions *A.1-A.6*, for a Nash equilibrium, the adversarial agent will choose $a_i^*\left(\varepsilon\right)$ defined as

$$a_i^*\left(\varepsilon\right)=\bar{a}_i+\varepsilon\frac{J_{a_i}\left(u^*\left(\varepsilon\right),a^*\left(\varepsilon\right)\right)}{\left(\sum_i J_{a_i}\left(u^*\left(\varepsilon\right),a^*\left(\varepsilon\right)\right)^2\right)^{1/2}}=\bar{a}_i+\varepsilon v_i\left(\varepsilon\right) \tag{5}$$

and the policymaker will choose

$$u^*\left(\varepsilon\right)=u^*\left(0\right)-\varepsilon J_{uu}\left(u^*\left(0\right),a^*\left(0\right)\right)^{-1}J_{ua}\left(u^*\left(0\right),a^*\left(0\right)\right)v\left(0\right)+o\left(\varepsilon\right) \tag{6}$$

where $v_i\left(\varepsilon\right)=\dfrac{J_{a_i}\left(u^*\left(\varepsilon\right),a^*\left(\varepsilon\right)\right)}{\left(\sum_i J_{a_i}\left(u^*\left(\varepsilon\right),a^*\left(\varepsilon\right)\right)^2\right)^{1/2}}$.

**Proof.**

We first calculate $a^*\left(\varepsilon\right)$. The constrained optimization problem for the adversarial agent is the maximization of

$$L=J\left(u,a\right)+\lambda\left(\varepsilon^2-\sum_i\left(a_i-\bar{a}_i\right)^2\right) \tag{7}$$

conditional on the choice of $u$ by the policymaker. The first order necessary conditions for a constrained maximum are

$$J_{a_i}\left(u,a^*\right)=2\lambda\left(a_i^*-\overline{a}_i\right)\ \forall i \tag{8}$$

and

$$\varepsilon^2 \geq \sum_i\left(a_i^*-\overline{a}_i\right)^2=\sum_i\left(\frac{J_{a_i}\left(u,a^*\right)}{2\lambda}\right)^2 . \tag{9}$$

The regularity conditions above ensure that (9) holds with equality at $a^*\left(\varepsilon\right)$ for $\varepsilon$ near 0. Combining eqs. (8) and (9), we have, using the second-order necessary condition to determine the sign in (10) below,

$$a_i^*\left(\varepsilon\right)=\overline{a}_i+\varepsilon\frac{J_{a_i}\left(u,a^*\right)}{\left(\sum_i J_{a_i}\left(u,a^*\right)^2\right)^{1/2}} \tag{10}$$

Eq. (10) immediately implies (5) using the definition for $v_i\left(\varepsilon\right)$ that appears in the Proposition. Note that the associated vector $v\left(\varepsilon\right)$(whose $i$'th element is $v_i\left(\varepsilon\right)$) is continuous in $\varepsilon$ at $\varepsilon=0$ by A1 and A6.

To calculate $u^*\left(\varepsilon\right)$, we start with the first order condition

$$J_u\left(u^*\left(\varepsilon\right),a^*\left(\varepsilon\right)\right)=0 \tag{11}$$

which holds as an immediate consequence of Conditions A4 and A5 Total differentiation of (11) with respect to $\varepsilon$ implies, when (11) is evaluated at $\varepsilon=0$,

$$J_{uu}\left(u^*(0),a^*(0)\right)\frac{du^*(0)}{d\varepsilon}+J_{ua}\left(u^*(0),a^*(0)\right)v(0)=0 \qquad (12)^4$$

so that

$$\frac{du^*(0)}{d\varepsilon}=-J_{uu}\left(u^*(0),a^*(0)\right)^{-1}J_{ua}\left(u^*(0),a^*(0)\right)v(0) \qquad (13)$$

which immediately implies the robust choice of control described by (6). Note that $\lim_{\varepsilon\to 0}\frac{o(\varepsilon)}{\varepsilon}=0$.

The formulas described by Proposition 1 are quite convenient as they illustrate how one can compute locally robust analogs of optimal controls using a set of original controls and various first and second derivatives of the loss function. These formulas generalize Brock, Durlauf and West (2003) which considered local robustness against scalar parameter uncertainty.

### Stackelberg equilibrium

An alternative to the Nash approach is the Stackelberg approach in which the optimizer is the leader. In this case, the choice of the parameter value by the adversarial agent may be modeled via a reaction function $a^*=R(u,\varepsilon)$. As before, the adversarial agent will exhaust the full set of constraints implied by eq. (2). The Stackelberg equilibrium thus requires that

$$J\left(u^*(\varepsilon),R\left(u^*(\varepsilon),\varepsilon\right)\right)\le J\left(u(\varepsilon),R\left(u(\varepsilon),\varepsilon\right)\right)\ \forall\ u\in R^n \qquad (14)$$

---

[4] The second term on the left hand side of eq. (12) follows from $\frac{d\left(\varepsilon v(\varepsilon)\right)}{d\varepsilon}=v(\varepsilon)+\varepsilon\frac{dv(\varepsilon)}{d\varepsilon}$ which equals $v(0)$ at $\varepsilon=0$.

and

$$R(u,\varepsilon) = \text{argmax}_{\text{a}}\ \{J(u,a) \text{ such that } \|a-\overline{a}\| \le \varepsilon\}. \tag{15}$$

It turns out that the Nash and Stackelberg equilibria are approximately equivalent, as stated in Proposition 2.

**Proposition 2. Approximate equivalence of Nash and Stackelberg decisions for robust policy problem**

Suppose that a policymaker chooses $u$ in order to minimize $J(u,a)$ whereas an adversarial agent chooses $a$ subject to (4) in order to maximize $J(u,a)$. Suppose that the policymaker is the leader in a Stackelberg game in which these choices are made and that the adversarial agent's reaction function is required to fulfill (15). Under assumptions *A.1-A.6*, for small enough $\varepsilon$, the Stackelberg decisions of the two agents will be within $o(\varepsilon)$ of the Nash decisions.

**Proof.**

By analogy to the previous analysis, it is straightforward to verify that

$$R_i(u,\varepsilon) = a_i^*(\varepsilon) = \overline{a}_i + \varepsilon \frac{J_{a_i}(u,a^*(\varepsilon))}{\left(\sum_i J_{a_i}(u,a^*(\varepsilon))^2\right)^{1/2}} = \overline{a}_i + \varepsilon v_i(u,\varepsilon). \tag{16}$$

The first order necessary condition associated with an optimum choice of $u$ implies

$$\begin{aligned} &\frac{dJ\left(u^*(\varepsilon), R\left(u^*(\varepsilon), \varepsilon\right)\right)}{du^*(\varepsilon)} \\ &= J_u\left(u^*(\varepsilon), R\left(u^*(\varepsilon), \varepsilon\right)\right) + \frac{dR\left(u^*(\varepsilon), \varepsilon\right)'}{du} J_a\left(u^*(\varepsilon), R\left(u^*(\varepsilon), \varepsilon\right)\right) \\ &= J_u\left(u^*(\varepsilon), R\left(u^*(\varepsilon), \varepsilon\right)\right) + \varepsilon v_u\left(u^*(\varepsilon), \varepsilon\right)' J_a\left(u^*(\varepsilon), R\left(u^*(\varepsilon), \varepsilon\right)\right) \\ &= 0 \end{aligned} \tag{17}$$

Total differentiation of (17) with respect to $\varepsilon$ implies

$$\begin{aligned} &J_{uu}\left(u^*(\varepsilon), R\left(u^*(\varepsilon), \varepsilon\right)\right)\frac{du^*(\varepsilon)}{d\varepsilon} + J_{ua}\left(u^*(\varepsilon), R\left(u^*(\varepsilon), \varepsilon\right)\right) R_\varepsilon\left(u^*(\varepsilon), \varepsilon\right) \\ &+ v_u\left(u^*(\varepsilon), \varepsilon\right)' J_a\left(u^*(\varepsilon), R\left(u^*(\varepsilon), \varepsilon\right)\right) + \varepsilon \frac{d\left(v_u\left(u^*(\varepsilon), \varepsilon\right)' J_a\left(u^*(\varepsilon), R\left(u^*(\varepsilon), \varepsilon\right)\right)\right)}{d\varepsilon} \\ &= 0 \end{aligned} \tag{18}$$

When $\varepsilon = 0$, (18) simplifies to

$$\begin{aligned} &J_{uu}\left(u^*(0), R\left(u^*(0), 0\right)\right)\frac{du^*(0)}{d\varepsilon} + J_{ua}\left(u^*(0), R\left(u^*(0), 0\right)\right) v\left(u^*(0), 0\right) \\ &+ v_u\left(u^*(0), 0\right)' J_a\left(u^*(0), R\left(u^*(0), 0\right)\right) = 0 \end{aligned} \tag{19}$$

Comparing (19) to (12), the only difference between the Nash solution and the Stackelberg solution is the additional term $v_u\left(u^*(0), 0\right)' J_a\left(u^*(0), R\left(u^*(0), 0\right)\right)$ in (19). The *m*'th term of this vector may be rewritten as

$$\sum_i J_{a_i}\left(u^*(0),R\left(u^*(0),0\right)\right)\frac{\partial\left(\frac{J_{a_i}\left(u^*(0),R\left(u^*(0),0\right)\right)}{\left(\sum_i J_{a_i}\left(u^*(0),R\left(u^*(0),0\right)\right)^2\right)^{1/2}}\right)}{\partial u_m} \tag{20}$$

To determine the value of (20), we employ

$$\sum_i\left(\frac{J_{a_i}\left(u^*(0),R\left(u^*(0),0\right)\right)}{\left(\sum_i J_{a_i}\left(u^*(0),R\left(u^*(0),0\right)\right)^2\right)^{1/2}}\right)^2\equiv 1 \tag{21}$$

Partial differentiation of (21) with respect to $u_m$ implies

$$\sum_i 2\frac{\partial\left(\frac{J_{a_i}\left(u^*(0),R\left(u^*(0),0\right)\right)}{\left(\sum_i J_{a_i}\left(u^*(0),R\left(u^*(0),0\right)\right)^2\right)^{1/2}}\right)}{\partial u_m}\left(\frac{J_{a_i}\left(u^*(0),R\left(u^*(0),0\right)\right)}{\left(\sum_i J_{a_i}\left(u^*(0),R\left(u^*(0),0\right)\right)^2\right)^{1/2}}\right).$$
$$=0$$

Since, *k* was arbitrary, comparison of (20) and (21) reveals that

$$J_a\left(u^*(0),R\left(u^*(0),0\right)\right)' v_u\left(u^*(0),0\right)=0 \tag{22}$$

Hence the Stackelberg solution coincides with the Nash solution up to terms that are of order $o(\varepsilon)$ which verifies Proposition 2. Note that one can perform analogous calculations and verify that the Nash equilibrium and Stackelberg equilibrium also coincide up to terms of order $o(\varepsilon)$ if the adversarial agent is the leader. Proposition 2 is useful as it shows that our local robustness analysis does not depend on a particular solution concept for the noncooperative game between the policymaker and the adversarial agent at least with respect to two important equilibrium notions for such models.

The equivalence of Nash and Stackelberg solutions appears in a number of contexts in game theory. For example, Hofbauer and Sorin (2002) locate a set of useful sufficient conditions for the value of a general zero-sum game to exist and, implicitly, for Nash and Stackelberg solutions to coincide. One may view our results as a local adaptation of work such as Hofbauer and Sorin.

### discrete alternatives

One can extend our analysis to consider environments in which there are discrete local alternatives to a given baseline model. This type of formulation seems appropriate for a number of macroeconomic contexts. For example, one can consider a baseline real business cycle model for which different types of nominal rigidities represent alternative directions along which to modify the baseline. More generally, formulating model uncertainty in terms of local discrete alternatives will make sense when there is prior information that delimits the directions along which model uncertainty exists. For each of the baseline models, one can perform a robustness analysis of the type we have described.

Once one considers discrete directions along which model uncertainty occurs, a new set of questions arise in terms of understanding the dispersions of payoffs and policies across the alternatives. Brock, Durlauf and West (2004) refer to these as outcome variance and action variance respectively, and provide a detailed discussion in the context of monetary policy rules. Here, we outline the notions of outcome and action variance for our general framework; where outcomes correspond to losses and actions correspond to choices of control. For simplicity, we treat $u$ as scalar.

Suppose that there exists a family of $q$ local alternatives to a given baseline model $J\left(u,\overline{a}\right)$ that are represented by

$$J\left(u,\overline{a}+\varepsilon l_q\right),\ q=1...Q. \tag{23}$$

where each $l_q$ is an $m\times 1$ vector that describes a direction along which the choice of the adversarial agent may vary the baseline Associated with each of these directions is an optimal control choice,

$$u^*\left(\overline{a}+\varepsilon l_q\right)=\text{argmax}_{\text{u}}\left(J\left(u,\overline{a}+\varepsilon l_q\right)\right) \tag{24}$$

and associated loss

$$J^*\left(\overline{a}+\varepsilon l_q\right)=J\left(u^*\left(\overline{a}+\varepsilon l_q\right),\overline{a}+\varepsilon l_q\right). \tag{25}$$

To compute the variances of these two quantities, first consider an expansion of (24) around $\varepsilon=0$,

$$u^*\left(\overline{a}+\varepsilon v_q\right)=u^*\left(\overline{a}\right)+u_a^*\left(\overline{a}\right)\varepsilon l_q+o\left(\varepsilon\right) \tag{26}$$

where

$$u_a^*\left(\overline{a}\right)=-J_{uu}\left(u^*\left(\overline{a}\right),\overline{a}\right)^{-1}J_{ua}\left(u^*\left(\overline{a}\right),\overline{a}\right) \tag{27}$$

by the first order conditions that determine $u^*\left(\overline{a}\right)$. By the envelope theorem,

$$J^*\left(\bar{a}+\varepsilon l_q\right) = J\left(u^*\left(\bar{a}+\varepsilon l_q\right), \bar{a}+\varepsilon l_q\right)$$
$$= J^*\left(\bar{a}\right) + \varepsilon J_a\left(u\left(\bar{a}\right), \bar{a}\right)' l_q + o\left(\varepsilon\right). \tag{28}$$

Equations (26) and (28) allow one to calculate the variance of controls and losses up to terms of $o\left(\varepsilon^2\right)$,

$$\operatorname{var}\left(u^*\left(\bar{a}+\varepsilon l_q\right)\right) =$$
$$\varepsilon^2\left(J_{uu}\left(u^*\left(\bar{a}\right),\bar{a}\right)^{-1} J_{ua}\left(u^*\left(\bar{a}\right),\bar{a}\right)\right)\operatorname{var}\left(l_q\right)\left(J_{uu}\left(u^*\left(\bar{a}\right),\bar{a}\right)^{-1} J_{ua}\left(u^*\left(\bar{a}\right),\bar{a}\right)\right)' + o\left(\varepsilon^2\right) \tag{29}$$

and

$$\operatorname{var}\left(J^*\left(\bar{a}+\varepsilon l_q\right)\right) = \varepsilon^2 J_a\left(u\left(\bar{a}\right),\bar{a}\right)' \operatorname{var}\left(l_q\right) J_a\left(u\left(\bar{a}\right),\bar{a}\right) + o\left(\varepsilon^2\right). \tag{30}$$

The formulas (29) and (30) provide a local analysis of the role of the loss function $J$ in transmitting underlying model uncertainty into two action variances and outcome variances. There are four main possibilities: action variance is small (large); outcome variance is small (large). Presumably, a policymaker prefers that both the action variance and the outcome variance are small; such preferences are often assumed in the monetary policy rules literature.

An important feature of eqs. (29) and (30) is the role that the loss function $J$ plays in contracting or expanding the "baseline" variance, $\operatorname{var}\left(l_q\right)$. Since the formulas that map the first and second derivatives of the loss function with respect to the vectors $a$ and *u*, the way the loss function does this involves the interactions of many terms.

## 3. Design limits and local robustness in discrete time systems

In this section, we consider the question of the design of locally robust controls for a scalar discrete time system. This sort of model frequently appears in macroeconomics contexts. Our analysis is restrictive in two respects: we focus on scalar state systems and we do not address issues of forward-looking elements in the state equation; we are pursuing the analysis of vector state systems and forward-looking models in subsequent work.

In this system, $x_t$ denotes the state of the system and $u_t$ denotes the control that is available to the policymaker. The law of motion for the state is

$$x_t = A(L)x_{t-1} + B(L)u_{t-1} + \xi_t \tag{31}$$

where the Wold representation of $\xi_t$ is denoted

$$\xi_t = w(L)\nu_t \, . \tag{32}$$

It is not necessary for the innovations $\nu_t$ to be independent or identically distributed; for simplicity we will assume that the innovations have a common variance. For analytical convenience, we assume that $w(L)$ is invertible; we will exploit the fact that in the Wold representation, one can always normalize the innovations $\nu_t$ so that $w_0 \equiv 1$, cf. Ash and Gardner (1975); we employ this normalization throughout.

We consider control rules of the form

$$u_{t-1} = -F(L)x_{t-1} \, . \tag{33}$$

This means that the state equation may be rewritten as

$$x_t = \left(A(L) - F(L)B(L)\right)x_{t-1} + w(L)\nu_t \, . \tag{34}$$

So long as (34) is asymptotically stable, $x_t$ will possess a spectral density. In general, it is possible for (34) to be unstable; we ignore this issue. Of course, an important objective in monetary policy may well be ensuring stability and so we implicitly confine ourselves to policies that do so.

We assume that the policymaker possesses a loss function

$$J = Ex^2, \tag{35}$$

which means that the sole objective of the policymaker is to stabilize the state variable. Our analysis may be generalized so that the policymaker also desires to minimize variability in the control, but we ignore this for expositional purposes.

## Design limits and optimal design

Our analysis of this environment will be facilitated by considering the frequency domain analogs of (31) to (34). We do this because the robustness issues that arise in this problem may best be understood as arising from uncertainty about the frequency-specific aspects concerning the law of motion for the state, as will become clear shortly. Frequency domain approaches to robust policy evaluation are also explored in Hansen and Sargent (2003a), Kasa (2000) and Sargent (1999).

Our approach follows Kwakernaak and Sivan (1972, ch. 6). Throughout, for any lag polynomial $C(L)$, the Fourier transform of its coefficients, $C\left(e^{-i\omega}\right) \equiv C(\omega)$, is defined as $\sum_{j=-\infty}^{\infty} C_j e^{-ij\omega}$. In analyzing this model, we work with the following transfer functions, which describe how shocks, state variables, and controls are interrelated:[5]

1. the transfer function that relates the control to the state,

---

[5] Transfer functions such as the ones we employ are valuable in that they allow one to identify the response of different aspects of the system (state, control, etc) to sinusoidal inputs and as such naturally arise in the frequency domain interpretation of the system.

$$H\left(\omega\right)=\left(e^{i\omega}-A\left(e^{-i\omega}\right)\right)^{-1}B\left(e^{-i\omega}\right), \tag{36}$$

2. the transfer function for the shocks to the state (when the control is set equal to zero)

$$h\left(\omega\right)=\left(e^{i\omega}-A\left(e^{-i\omega}\right)\right)^{-1}, \tag{37}$$

and

3. the transfer function from the state to the control

$$G\left(\omega\right)=F\left(e^{-i\omega}\right). \tag{38}$$

In order to identify the effects of the control rule on the variance of the state, it is useful to express the variance in terms of frequency-specific contributions. It is straightforward to show that[6]

$$Ex_t^2=\int_{-\pi}^{\pi}\frac{f_{\xi}\left(\omega\right)}{\left|h\left(\omega\right)^{-1}+B\left(\omega\right)G\left(\omega\right)\right|^2}d\omega \tag{39}$$

which employs the $h\left(\omega\right)$ and $G\left(\omega\right)$ functions we have defined.[7] This expression may be rewritten as

$$Ex_t^2=\int_{-\pi}^{\pi}\left|S\left(\omega\right)\right|^2 f_{x^{NC}}\left(\omega\right)d\omega \tag{40}$$

---

[6] Eq. (39) is an immediate implication of Kwakernaak and Sivan (1972), Theorem 6.21, pg. 469.

[7] Note that $f_{\xi}\left(\omega\right)=\frac{\sigma_v^2}{2\pi}\left|w\left(e^{-i\omega}\right)\right|^2$

where $S(\omega)$ is called the sensitivity function and is defined by

$$S(\omega)=\frac{1}{1+H(\omega)G(\omega)} \tag{41}$$

and $x_t^{NC}$ is defined as the $x_t$ process under the counterfactual that the control is set equal to zero each period, i.e.

$$x_t^{NC}=A(L)x_{t-1}^{NC}+\xi_t \tag{42}$$

with associated spectral density

$$f_{x^{NC}}(\omega)=\left|h(\omega)\right|^2 f_\xi(\omega). \tag{43}$$

Notice that $f_{x^{NC}}(\omega)$ does not depend on the control. The sensitivity function does depend on the form of the control via its dependence on $H(\omega)$ and $G(\omega)$. These transfer functions reflect the feedback from the control to the state and state to the control respectively, the sensitivity function thus allows for explicit calculation of how a given control, working through these two transfer functions, affects the frequency-specific variance components of $f_x(\omega)$.

Equation (40) provides a frequency domain perspective on optimal policy design. Each possible choice of a control rule $F(L)$ determines an associated sensitivity function $S(\omega)$, as indicated by eqs. (38) and (41). Hence the optimal control problem is equivalent to finding, among the set of feasible $S(\omega)$ functions the one that minimizes (40). From the perspective of (40), a control can perfectly stabilize the state if and only if it is possible that $\left|S(\omega)\right|^2=0\ \forall\omega$. In general, this is impossible since the control rule does not allow conditioning on $\nu_t$. Given this restriction, at first glance, it would seem that minimizing (40) subject to (41) is achieved by setting $F(L)B(L)=A(L)$, so that the $x_t$ process is reduced to the error term $\nu_t$. In fact, this is the optimal solution if $\nu_t$ is uncorrelated. Our interest is in the general case where $\nu_t$ is correlated.

What restrictions does the structure of this problem impose on the choice of $S(\omega)$? There exists a general constraint on the choice of $S(\omega)$ known as the Bode integral constraint. The constraint states that

$$\int_{-\pi}^{\pi} \ln\left(\left|S(\omega)\right|^2\right) d\omega = K \tag{44}$$

where $K$ is a constant that depends on various aspects of eqs. (31)-(32); its exact value is defined in the Appendix. This constraint is critical in understanding optimal policy analysis as it characterizes fundamental limits in policy design. For our purposes, the important property of $K$ is that it must be nonnegative. This implies that it is impossible for $S(\omega)$ to always be less than 1. Hence, *any* choice of controls will increase the variance associated with at least one set of frequencies relative to what would occur were the control not to be implemented. This represents a fundamental tradeoff in all control problems of this type.

From the frequency domain perspective, the evaluation of the effects of a policy rule may be based on a comparison of the variance of the state produced with the rule versus the variance of the state when the control is set equal to zero in all periods. This difference is:

$$\begin{gathered}\int_{-\pi}^{\pi} f_{x^{NC}}(\omega)\, d\omega - \int_{-\pi}^{\pi} \left|S(\omega)\right|^2 f_{x^{NC}}(\omega)\, d\omega = \\ \int_{-\pi}^{\pi} \left(1-\left|S(\omega)\right|^2\right) f_{x^{NC}}(\omega)\, d\omega\end{gathered} \tag{45}$$

A policymaker will want to shape $S(\omega)$ so as to maximize (45). This indicates how one may construct optimal policy rules. The Bode integral constraint establishes the constraints that must be placed on this maximization problem. The optimal policy may therefore be conceptualized as the following constrained optimization problem, in which the policymaker chooses a sensitivity function to minimize the variance of the state subject to the Bode integral constraint (44):

$$\begin{gathered}\int_{-\pi}^{\pi} p(\omega)\, f_{x^{NC}}(\omega) d\omega + \lambda\left(K - \int_{-\pi}^{\pi} \ln p(\omega)\, d\omega\right) = \\ \lambda K + \int_{-\pi}^{\pi}\left( p(\omega)\, f_{x^{NC}}(\omega) - \lambda \ln p(\omega)\right) d\omega\end{gathered} \tag{46}$$

The first order necessary conditions for minimization imply that

$$f_{x^{NC}}(\omega) = \frac{\lambda}{p(\omega)} \rightarrow f_{x^{NC}}(\omega)\, p(\omega) = \lambda \tag{47}$$

This implies that the Fourier transform of the optimal feedback rule is implicitly defined by

$$p(\omega)\, f_{x^{NC}}(\omega) = \frac{\left|w\left(e^{-i\omega}\right)\right|^2 \sigma_v^2}{\left|e^{i\omega} - A\left(e^{-i\omega}\right) + B\left(e^{-i\omega}\right) F^*\left(e^{-i\omega}\right)\right|^2} = \lambda \tag{48}$$

The optimal feedback rule means that the $x_t$ process is white noise. This makes intuitive sense as the policymaker will use the control to eliminate any (linearly) predictable component of the state. This suggests that the control rule $u_t = -F^*(L)\, x_{t-1}$ should have the property that when the rule is implemented, $x_t = v_t$. By (34), for any control rule $u_{t-1} = -F(L)\, x_{t-1}$

$$\begin{gathered}x_t = A(L)\, x_{t-1} - B(L)\, F(L)\, x_{t-1} + w(L)\, v_t = \\ \left(A(L) L - B(L)\, F(L)\, L\right) x_t + w(L)\, v_t\end{gathered} \tag{49}$$

so that

$$\left(1-A(L)L+B(L)F(L)L\right)x_t = w(L)v_t \tag{50}$$

Therefore, the condition $x_t = v_t$ requires that the optimal feedback rule $-F^*(L)$ is consistent with

$$1-A(L)L+B(L)F^*(L)L = w(L) \tag{51}$$

This argument is heuristic in that it is based on first-order conditions and the class of lag operators which span the space of potential policy rules is not precisely defined. One can formally prove that the implied optimal policy rule in (51) is in fact optimal.

**Proposition 3. Optimal Policy Rule**

Assume that the lag operators $A(L)$, $B(L)$ and $w(L)$ are all polynomials of finite degree. The optimal feedback rule to minimize (35) given (31) is

$$-F^*(L) = B(L)^{-1}\left(L^{-1}-w(L)L^{-1}-A(L)\right) = B(L)^{-1}\left(-\left(w(L)L^{-1}\right)_+ - A(L)\right) \quad (52)^8$$

**Proof.**

We initially consider the case where $\xi_t$ is white noise, i.e. $w(L)=1$. In this case, defining $C(L)=A(L)-B(L)F(L)$ one can rewrite the state equation as

$$\left(1-LC(L)\right)x_t = \prod_k \left(1-r_kL\right)x_t = \xi_t \tag{53}$$

---

[8] For any lag polynomial $C(L)$, $C(L)_+$ denotes the part of the polynomial where all terms with negative exponents are dropped; the annihilation operator "+" is discussed in Sargent (1987), pg. 292, Notice that this last expression in (52) exploits the fact $w_0 = 1$ in our normalization of the Wold decomposition.

where the terms $r_k$ denote the roots of the finite degree lag polynomial on the right hand side of (53). Eq. (53) allows one to express the unconditional variance of the state as

$$Ex_t^2 = \frac{\sigma_\xi^2}{2\pi}\int_{-\pi}^{\pi}\frac{d\omega}{\left|e^{i\omega} - C\left(e^{-i\omega}\right)\right|^2} = \frac{\sigma_\xi^2}{2\pi}\int_{-\pi}^{\pi}\frac{d\omega}{\prod_k \left|1 - r_k e^{-i\omega}\right|^2} = \frac{\sigma_\xi^2}{2\pi}\int_{-\pi}^{\pi}\frac{d\omega}{\prod_k \left|e^{i\omega} - r_k\right|^2} \qquad (54)$$

The left hand side of (54) is minimized if its log is minimized. We minimize over all $r_k$ restricted to lie inside the unit circle in the complex plane. By Jensen's inequality, (54) implies

$$\begin{aligned} &\ln\sigma_\xi^2 + \ln\left(\frac{\sigma_\xi^2}{2\pi}\int_{-\pi}^{\pi}\frac{1}{\prod_k \left|1 - r_k e^{-i\omega}\right|^2}d\omega\right) \geq \\ &\ln\sigma_\xi^2 + \frac{1}{2\pi}\int_{-\pi}^{\pi}\ln\left(\frac{1}{\prod_k \left|1 - r_k e^{-i\omega}\right|^2}\right)d\omega = \ln\sigma_\xi^2 - \left(\frac{1}{\pi}\int_{-\pi}^{\pi}\ln\prod_k\left|1 - r_k e^{-i\omega}\right|d\omega\right) \end{aligned} \qquad (55)$$

Lemma 5 of Wu and Jonckheere (1992) shows that $\int_{-\pi}^{\pi}\ln\left|e^{i\omega} - r_k\right|d\omega = 0$ if $\left|r_k\right| \leq 1$. Hence, the right hand side of (55) is $\ln\sigma_\xi^2$. This gives us a lower bound on the objective function; hence a choice of $r_k$ (via the choice of $F(L)$) that achieves this bound must be optimal. The choice, $r_k = 0,\ \forall k$ achieves this bound. By (54), $r_k = 0,\ \forall k$ means that $C(L) = 0$, hence

$$-F^*(L) = \frac{A(L)}{B(L)} \qquad (56)$$

which corresponds to (52) when $w(L) = 1$.

For general invertible $w(L)$ with $w_0 = 1$,

$$Ex_t^2 = \frac{\sigma_\xi^2}{2\pi}\int_{-\pi}^{\pi} \frac{\left|w\left(e^{-i\omega}\right)\right|^2}{\prod_k \left|e^{i\omega} - C\left(e^{-i\omega}\right)\right|^2} d\omega = \frac{\sigma_\xi^2}{2\pi}\int_{-\pi}^{\pi} \frac{\prod_l \left|e^{i\omega} - \mu_l\right|}{\prod_k \left|e^{i\omega} - r_k\right|^2} d\omega \tag{57}$$

The same argument as above implies that the lower bound on $\ln Ex_t^2$ is $\ln \sigma_\nu^2$. This bound is achieved if

$$\frac{1}{2\pi}\int_{-\pi}^{\pi} \frac{\left|w\left(e^{-i\omega}\right)\right|^2}{\prod_k \left|e^{i\omega} - C^*\left(e^{-i\omega}\right)\right|^2} d\omega = 1 \tag{58}$$

i.e.

$$1 - LC^*\left(L\right) = w\left(L\right) \tag{59}$$

If we expand out (59), using $w_0 = 1$,

$$1 - L\left(c_o + c_1 L + c_2 L^2 \ldots\right) = 1 + w_1 L + w_2 L^2 \ldots \tag{60}$$

which makes clear that $C\left(L\right) = A\left(L\right) - B\left(L\right)F\left(L\right)$ is a one-sided polynomial in $L$, completes the proof.

This analysis indicates a general strategy for optimal policy construction when the model is known. Intuitively, a policymaker chooses a feedback rule to eliminate any predictability in the state. Further, the calculations involved in the construction of an optimal policy in the absence of model uncertainty allow one to develop robust analogs of the optimal policy rule when model uncertainty is present.

Our minimization of $Ex_t^2$ over a class of feedback rules is closely related to optimization over a set of stationary states in general capital theory (Marimon, 1989; McKenzie, 1986) as well as Woodford's notion of "optimal from a timeless perspective" (Woodford, 2003a, Section 3, Chapter 7 and Section 1.1 Chapter 8). In the notion of optimization over steady states in Marimon (1989) and McKenzie (1986) one optimizes over a set of stationary strategies so that transient effects from historically-given initial conditions are ignored. Hence time inconsistencies that arise from discretionary optimization from initial conditions that are "reset" every period do not arise. Since this notion of a steady state is timeless there is a direct connection with Woodford's concept. In taking this approach, it is important to remember that the problem of time inconsistency must be dealt with in our framework once one moves beyond the class of stationary policies considered here.[9]

**design limits and robustness**

Our analysis of design limits provides a framework for understanding robustness via the role of the Bode integral constraint in determining optimal policy. For the control problem we have just described, suppose that model uncertainty relates to the spectral density of the innovation process $\xi_t$, i.e. $f_\xi(\omega)$. Uncertainty about $f_\xi(\omega)$ is a flexible way of allowing for uncertainty about the process that describes the state, eq. (31). Substantively, our assumption to restrict model uncertainty to uncertainty about $f_\xi(\omega)$ implies two restrictions on the policymaker's information set: first, the policymaker knows the autoregressive structure polynomial $A(L)$, and second, the policymaker knows the feedback parameter from the policy rule to the state, $B(L)$. Notice that this prior knowledge still allows for rich forms of model uncertainty as the policymaker does not know the impulse response function associated with the innovations $\nu_t$. Further, uncertainty about $A(L)$ may typically be translated into uncertainty about $w(L)$. At least metaphorically, the sorts of model uncertainty we allow corresponds to Milton

---

[9] Further, in more elaborate frameworks, such as those with multiple variables to stabilize with a given control or various forward looking models, it may not be possible or even optimal to make each series unpredictable. An example of a model where optimal rules do not produce white noise series is Woodford (2003b). We thank a referee for these observations.

Friedman's famous concern (1948) about long and variable lags in the mapping from feedback rules to macroeconomic aggregates.

The implications of this type of model uncertainty can be analyzed using the results in Section 2. We focus on the Nash equilibrium between the policymaker and adversarial agent; as shown in Section 2, the Stackelberg equilibrium will produce identical results up to order $\varepsilon$. Relative to the $J(u,a)$ formalism of Section 2, $p(\omega)$ plays the role of the control vector $u$ and the function $f_\xi(\omega)$ plays the role of the parameter vector $a$. The applicability of the results in Section 2 is not affected by our use of functions rather than vectors for local robustness analysis. At an abstract level, an $n$-dimensional vector $u$ is a map from the set $\{1,2,....N\}$ to the real line (i.e. each component maps to a value on the real line) just as a function is a map from some set $S$ to the real line. The derivation of first order necessary conditions for Nash equilibrium for functions $p(\omega)$ and $f_\xi(\omega)$ follows the same logic as was employed for the vectors $u$ and $a$, i.e. one employs a variational analysis component by component (i.e. for each value of the domain $S$) to derive the first order conditions for the best response of each side of the noncooperative game. While there are some technical details in the case of functions with a continuous number of "components" in contrast to vectors which have only a finite number of components, from the perspective of robustness analysis one proceeds exactly as in the vector case of doing calculus with respect to each component, i.e. each argument, of the function. Hence one can proceed as in standard calculus and get the right answer. See any textbook in the calculus of variations such as Kamien and Schwartz (1991) for this basic method of deriving first order necessary conditions. Indeed our case is much easier because there are no dynamics that have to be considered unlike the general case treated in Kamien and Schwartz (1991).

Following our earlier notion we start with a baseline spectral density $\overline{f}_\xi(\omega)$ and define model uncertainty as meaning that the true spectral density lies in a set defined around this baseline, in other words, the set of potential models is defined by all spectral densities $f_\xi(\omega)$ such that

$$\int_{-\pi}^{\pi} \left( f_\xi(\omega) - \overline{f}_\xi(\omega) \right)^2 d\omega \leq \varepsilon^2 . \tag{61}$$

This approach to model uncertainty is closely related to that in Sargent (1999).[10] While this approach assumes that the policymaker knows $B(L)$ and $A(L)$, it is in fact relatively general since, from the perspective of $x_t^{NC}$, fairly general uncertainty about the spectral density $f_{x^{NC}}(\omega)$ can be generated via uncertainty in $w(L)$, which as noted in footnote 7 determines the spectral density $f_\xi(\omega)$.

From this model space, the adversarial agent chooses a spectral density in order to maximize

$$J\left( p^*(\omega), f_\xi(\omega) \right) = Ex_t^2 \tag{62}$$

conditional on the policymaker's choice $p^*(\omega)$ while the policymaker chooses a policy rule that robustifies against the adversarial agent's behavior. We now calculate these behaviors.

For the adversarial agent, the optimal choice $f_\xi^*(\omega)$ may be computed by applying the general formula (5) to the model space defined by eq. (61). This requires calculating the version of $J_{a_i}$ that applies to the policy problem we have described. For notational consistency, we will define $a_i = f_\xi(\omega)$ so that the marginal change in loss function for a frequency-specific change in the spectral density is $J_{f_\xi(\omega)}$. This partial derivative can be calculated using (40) and (47). The marginal effect of a change in the value of $\overline{f}_\xi(\omega)$ at a given frequency on $Ex_t^2$, is $p^*(\omega)\left|h(\omega)\right|^2$; this follows immediately

---

[10] Comparing our eq. (61) and Sargent (1999 eq. (5)), there are two differences. First, we do not require all spectral densities in the model space to integrate to the same unconditional variance for model errors $\xi_t$. Second, we do not assume that the baseline spectral density is that of a white noise process.

from the frequency definition of the variance and the envelope theorem. Since the policymaker's loss function is $Ex_t^2$, this implies that

$$J_{f_\xi(\omega)} = p^*(\omega)\left|h(\omega)\right|^2 . \tag{63}$$

Eq. (63) shows that $J_{f_\xi(\omega)} > 0 \ \forall \omega$, which is necessary to allow one to apply (5) to this case.[11] By (47), $p^*(\omega)$, the Nash equilibrium policymaker's choice at the baseline spectral density $\overline{f}_\xi(\omega)$, must fulfill $p^*(\omega) = \lambda \overline{f}_{x^{NC}}(\omega)^{-1}$. By (47), we can rewrite this expression as

$$p^*(\omega) = \frac{\lambda \overline{f}_\xi(\omega)^{-1}}{\left|h(\omega)\right|^2} . \tag{64}$$

Substituting (63) and (64) into (5), the adversarial agent chooses a spectral density of the form

$$f_\xi^*(\omega) = \overline{f}_\xi(\omega) + \varepsilon \frac{\overline{f}_\xi(\omega)^{-1}}{\left\| \overline{f}_\xi(\omega)^{-1} \right\|} + o(\varepsilon) . \tag{65}$$

This formula reveals an important feature of how an adversarial agent will allocate spectral power relative to a baseline spectral density. The expression $\frac{\overline{f}_\xi(\omega)^{-1}}{\left\| \overline{f}_\xi(\omega)^{-1} \right\|}$

---

[11] It is important to verify that Assumption A.3 holds for models of this type as there are some forms of model uncertainty where the assumption fails. For example, if the state equation is $x_t = A_0 x_{t-1} + u_{t-1} + \xi_t$, with $\xi_t$ white noise, if model uncertainty occurs because $A_0$ is not known, i.e. it is only known that it lies in the interval $\left[\overline{A}_0 - \varepsilon, \overline{A}_0 + \varepsilon\right]$, the model will not fulfill A.3 and in fact there will not exist a Nash equilibrium in pure strategies.

is large when $\overline{f}_\xi(\omega)^{-1}$ is large, i.e. when $\overline{f}_\xi(\omega)$ is small. Hence, the adversarial agent will increase the relative power at frequencies that contribute relatively little variance at the baseline spectral density. The result may be understood intuitively using the Bode integral constraint. A policymaker, because of this constraint, originally chooses a feedback rule such that $\left|S(\omega)\right|$ is small at those frequencies at which $\overline{f}_\xi(\omega)$ is large. Hence, the most effective allocation of spectral power by the adversarial agent is to those frequencies whose contributions to the variance of the state have not been downweighted by the original feedback rule chosen by the policymaker.

To derive a sense of what eq. (65) means operationally, we consider the case where the baseline process for $\xi_t$ is AR(1), i.e. $\xi_t = .25\xi_{t-1} + \nu_t$. The associated baseline spectral density is, for $\sigma_\nu^2 = 1$,

$$\overline{f}_\xi(\omega) = \frac{1}{2\pi\left(1 - .5\cos\omega + .0625\right)} \tag{66}$$

Figure 1 provides the baseline as well as the adversarial agent's choice of $f_\xi^*(\omega)$. As the Figure illustrates, while the adversarial agent adds power to each frequency, the additions are smallest at the low frequencies and largest at the high frequencies, which is of course what is described by (65).

How should a policymaker with an aversion to ambiguity respond to model uncertainty about $f_\xi(\omega)$? This question may be addressed immediately by employing eq. (47) under the assumption that $f_\xi(\omega) = f_\xi^*(\omega)$. From (47), it must be the case that

$$p^*(\omega) \propto \left|h(\omega)\right|^2 f_\xi^*(\omega) \tag{67}$$

where the constant of proportionality is implicitly determined by the Bode integral constraint. The robust solution can be quite complex as the adversarial agent's choice of $f_\xi^*(\omega)$ can convert the AR(1) process into a complicated ARMA structure.

**simple policy rules: an example**

We conclude this section with the analysis of an example of how to use these formulas to study robustness in the context of simple rules.[12] Suppose that $A(L)=0$ and $w(L)=1+\theta L$, $|\theta|<1$. In this case, $x_t^{NC}$ is an MA(1) process. Suppose that a policymaker is constrained to follow rules of the form:

$$u_t = -fx_{t-1}. \tag{68}$$

Following our earlier notation, $f$ corresponds to $u$ and $\theta$ corresponds to *a*, so that this model has an associated loss function

$$J(u,a) = \frac{\sigma_\nu^2}{2\pi}\int_{-\pi}^{\pi}\frac{\left|1+\theta e^{-i\omega}\right|^2}{\left|1-ue^{-i\omega}\right|^2}d\omega = \frac{\sigma_\nu^2}{2\pi}\int_{-\pi}^{\pi}\frac{\left|e^{i\omega}+\theta\right|^2}{\left|e^{i\omega}-u\right|^2}d\omega = \frac{\sigma_\nu^2}{2\pi}\int_{-\pi}^{\pi}\frac{a(\omega)}{\left|e^{i\omega}-u\right|^2}d\omega. \tag{69}$$

The substitution $a(\omega)\equiv\left|e^{i\omega}+\theta\right|^2$ in the last equality in (69) appears in order to make clear how the adversarial agent can affect the loss function frequency by frequency.

By Proposition 3 above, when there is no model uncertainty, the optimal choice of *f*, $f^*$, is defined by $f^*=-\theta$. It is straightforward to verify that this solution satisfies the second order necessary condition for minimizing $J(u,a)$ given our assumption that $|\theta|<1$. Relative to our abstract formulation, one can interpret $\bar{a}(\omega)=\left|e^{i\omega}+\theta\right|^2$ as the object chosen by the adversarial agent within a local space.

In order to calculate how local model uncertainty affects the choice of feedback parameter *f*, we employ eqs. (5) and (6), as before we determine the implicit behaviors

---

[12] This example partly builds on Brock and Durlauf (2004)'s discussion of design limits and non-time separable preferences. Unfortunately, the printed version of this paper contains some typographical errors; corrected text is available from the authors on request.

of the adversarial agent and policymaker frequency by frequency. Using (69), one can immediately calculate

$$J_{a(\omega)} = \frac{\sigma_v^2}{2\pi \left| e^{i\omega} - u \right|^2} \tag{70}$$

so that, using eq. (5),

$$a^*(\omega) = \bar{a}(\omega) + \varepsilon v(\omega) \tag{71}$$

where

$$v(\omega) = \frac{\dfrac{1}{\left| e^{i\omega} - u \right|^2}}{\left| \displaystyle\int_{-\pi}^{\pi} \frac{d\omega}{\left| e^{i\omega} - u \right|^2} \right|}. \tag{72}$$

The policymaker’s response is described by eq. (6) and may be written in our context as

$$\frac{du^*(0)}{d\varepsilon} = -\frac{1}{J_{uu}} \int_{-\pi}^{\pi} J_{ua(\omega)} v(\omega) \, d\omega. \tag{73}$$

The term $J_{uu}$ is positive for $|\theta| < 1$, so the sign of the change is determined by the integral part of (73). Differentiating (69) twice,

$$J_{ua(\omega)} = -\frac{\sigma_v^2 (u - \cos\omega)}{\pi \left| e^{i\omega} - u \right|^4}. \tag{74}$$

Combining (72) and (74), some algebraic manipulation yields

$$J_{ua}(\omega)v(\omega) \propto \frac{u-\cos\omega}{1+u^2-2u\cos\omega} = \frac{\theta+\cos\omega}{1+\theta^2+2\theta\cos\omega} \tag{75}$$

where the equality in (75) reflects $u^* = -\theta$. Therefore

$$\frac{du^*(0)}{d\varepsilon} \propto \int_{-\pi}^{\pi} \frac{\theta+\cos\omega}{1+\theta^2+2\theta\cos\omega} d\omega = 2\int_{0}^{\pi} \frac{\theta+\cos\omega}{1+\theta^2+2\theta\cos\omega} d\omega \tag{76}$$

This last integral equals 0, by Formula 2.554.2 of Gradshteyn and Ryzhik (2000). Therefore for this particular model, the introduction of local model uncertainty has no first order effect on the robustifying policymaker. While this result may not seem too exciting, what matters for our purposes is that the finding is far from obvious and our general robustness formulas provide a way of doing the calculations necessary to find this out. The result also indicates that just as robust policies can be more or less aggressive than policies chosen in the absence of model uncertainty (Brock, Durlauf, and West, 2003; Giannoni, 2002) it is also possible for the policies to be equivalent in a nonpathological case.

## 4. Some implications for monetary policy

In this section we illustrate some of the implications of the methods we have derived for monetary policy. We do this in the context of a one-dimensional version of the Rudebusch-Svensson model developed in Svensson (1996). In this model, the output gap, $y_t$, is the control and inflation, $\pi_t$, is the state variable. Inflation evolves according to

$$\pi_t = (1-\overline{\eta})\pi_{t-1} + by_{t-1} + \xi_t, \ \overline{\eta} \in (0,1) \tag{77}$$

and the policymaker is assumed to wish to minimize the loss function $E\pi_t^2$.

In the popular Taylor rule framework, monetary policy analysis is typically conducted in the context of simple rules. For this one-dimensional model, these rules may be parameterized by

$$y_{t-1} = \rho_y y_{t-2} + \rho_\pi \pi_{t-1} . \tag{78}$$

Without loss of generality, we assume that $\xi_t$ is uncorrelated and we normalize $b = 1$.

At the baseline parameter, $1-\overline{\eta}$, the policymaker will choose the control rule $y_{t-1} = -\left(1-\overline{\eta}\right)\pi_{t-1}$.[13] We consider the question of developing robust versions of the optimal monetary policy rules when we index models by the lagged inflation coefficient, $1-\overline{\eta}$. We therefore consider a set of possible models where the $q$'th model is associated with lagged inflation parameter $1-\overline{\eta}+\eta v_q$.

For each model $q$, we can calculate the (model-specific) Bode constant $K_q$ as

$$K_q = 4\pi \ln\left(1-\overline{\eta}+\eta v_q\right) \text{ if } \eta v_q > \overline{\eta},\ 0 \text{ otherwise }.^{14} \tag{79}$$

Therefore the Bode constant for model $q$ is positive if and only if $\eta v_q > \overline{\eta}$.

One can also calculate the $S\left(\omega\right)$ function and show that

$$\left|S\left(\omega\right)\right|^2 = \frac{\left|e^{i\omega} - \left(1-\overline{\eta}+\eta v_q\right)\right|^2}{\left|e^{i\omega} - \eta v_q\right|^2} . \tag{80}$$

Eq. (80) in turn implies $\left|S\left(\omega\right)\right|^2 > 1$ if and only if

---

[13] See Svensson (1996) for the analysis of optimal feedback rules which assume the form $y_{t-1} = \rho_\pi \pi_{t-1}$ for this model.

[14] This follows immediately from eq. (89) in the Appendix since $L\left(\omega\right) = \dfrac{1-\overline{\eta}}{e^{i\omega} - \left(1-\overline{\eta}+\eta v_q\right)}$.

$$1-\overline{\eta}+2\eta v_q > 2\cos\omega . \tag{81}$$

How may these calculations be used to inform a policymaker? Recall that the function $S(\omega)$ measures the sensitivity of the state to shocks of frequency $\omega$. A policymaker will want to concentrate on designing monetary policy rules which work especially well for those models where the Bode constraint is positive because these models will have relatively large sets of frequencies such that $\left|S(\omega)\right|^2 > 1$; such frequencies are precisely those whose variance contributions to the state will be magnified by an adversarial agent.

This case can also allow one to discuss robustness within and across models: a policymaker may be concerned about robustification against local misspecification of the exogenous shocks for *each* of these $Q$ separate models. Put differently, a local "cloud" of uncertainty about the spectral density of shocks may exist for each model. When there is also uncertainty about which of the $Q$ models is the true one, then there exist *two* layers of model uncertainty in this case. The Bode integral constraint can help provide insight into which of the $Q$ models has the most serious vulnerability to local misspecification of the spectral density of the innovation process and at which frequencies. In our view, this is an important area for future research, as it would allow for a combination of "global" and local model uncertainty, since there is no requirement that the $Q$ models are near each other in the way that models with different innovation processes were near each other as described by eq. (61).

Bode's constraint on $p(\omega)$ tells us that there is a frequency band where $\left|S(\omega)\right| > 1$ and this band will tend to be larger when the Bode constant, $K_q$, is strictly positive. Further, the Bode theory tells us that the frequencies most vulnerable to misspecifications for the primary agent are those frequencies where $p(\omega)\left|h(\omega)\right|^2$ is relatively large and, modulo the effect of $\left|h(\omega)\right|^2$, those frequencies are the ones where $p(\omega)$ is relatively larger. The Bode constraint warns the policymaker that frequencies exist where $\left|S\left(\omega\right)\right|>1$, i.e. misspecifications are magnified at those particular frequencies and it warns the primary agent to be particularly wary of local misspecification of the outside shocks when the Bode constant, $K_q$, is positive. The integration of the Bode constraint with robustness analysis suggests that a policymaker will be particularly concerned to robustify policies against misspecification of the outside shocks when there are possible baseline models with positive Bode constants.

These basic ideas may be used to begin to think about dynamic issues in policy analysis. Here we briefly sketch a sequence of interactions between the policymaker and the adversarial agent. Suppose that the policymaker and the adversarial agent each follow simple learning rules in that their decisions at time $t$ depend on their opponent's behavior at time $t-1$. In other words,

$$f_{\xi,t}^{*}\left(\omega\right)=\arg\max\ J\left(f_{\xi,t}\left(\omega\right),p_{t-1}\left(\omega\right)\right) \tag{82}$$

and

$$p_{t}^{*}\left(\omega\right)=\arg\min\ J\left(f_{\xi,t-1}\left(\omega\right),p_{t}\left(\omega\right)\right) \tag{83}$$

where argmax is taken over the set of spectral densities that satisfy (61). This type of model will produce a sequence of spectral density choices by the adversarial agent; following our earlier analysis, the sequence will obey.

$$f_{\xi,t}^{*}\left(\omega\right)=f_{\xi,t-1}^{*}\left(\omega\right)+\varepsilon\frac{p_{t-1}^{*}\left(\omega\right)\left|h\left(\omega\right)\right|^{2}}{\left\|p_{t-1}^{*}\left(\omega\right)\left|h\left(\omega\right)\right|^{2}\right\|}. \tag{84}$$

In response the sequence of control functions chosen by the policymaker must be such that

$$p_t^*(\omega) \propto \left|h(\omega)\right|^2 \left( f_{\xi,t-2}^*(\omega) + \varepsilon \frac{p_{t-2}^*(\omega)\left|h(\omega)\right|^2}{\left\| p_{t-2}^*(\omega)\left|h(\omega)\right|^2 \right\|} \right). \qquad (85)$$

Equations (82)-(85) suggest that interesting issues of policy robustness will arise in dynamic contexts, as the least favorable model for the policymaker will evolve across time in response to the evolving decisions of the adversarial agent. We leave formal analysis of this question to future work.

## 5. Summary and conclusions

This paper has provided an initial outline of a theory of local robustness. We have provided some general conditions that describe how, in the presence of local model uncertainty, robustness policy rules may be analyzed using basic calculus tools. This analysis provides explicit characterizations of the robust analog to an optimal control solution and the implied least favorable model for which the robust solution is an optimal control. We apply this basic framework to a discrete time scalar univariate optimal control problem. This necessitated development of a theory of design limits in such problems. Our analysis indicated how robustness considerations imply that a policymaker will, for certain types of model uncertainty around a baseline, guard against models of the shocks to a system that are less persistent than the baseline model. Finally, we have commented on the implications of this for monetary policy.

Our analysis provides an initial outline of tools for analyzing local robustness and integrating robustness into the analysis of limits to optimal policies. As such, many questions are left unaddressed. For example, we have not addressed the question of how to translate our abstract results on local robustness to dynamic models in which the loss function of the policymaker contains more than one argument or where the policymaker possesses more than one policy tool. Further, we have not addressed the important question of how learning by the policymaker may influence the way in which model uncertainty is conceptualized. Finally, we have not addressed how design limit results are affected by the presence of forward looking elements in the law of motion for the states of interest. Our intent in this paper is to illustrate some basic ideas that we believe are important in developing a general theory of policy design that recognizes the limitations of current macroeconomic understanding.

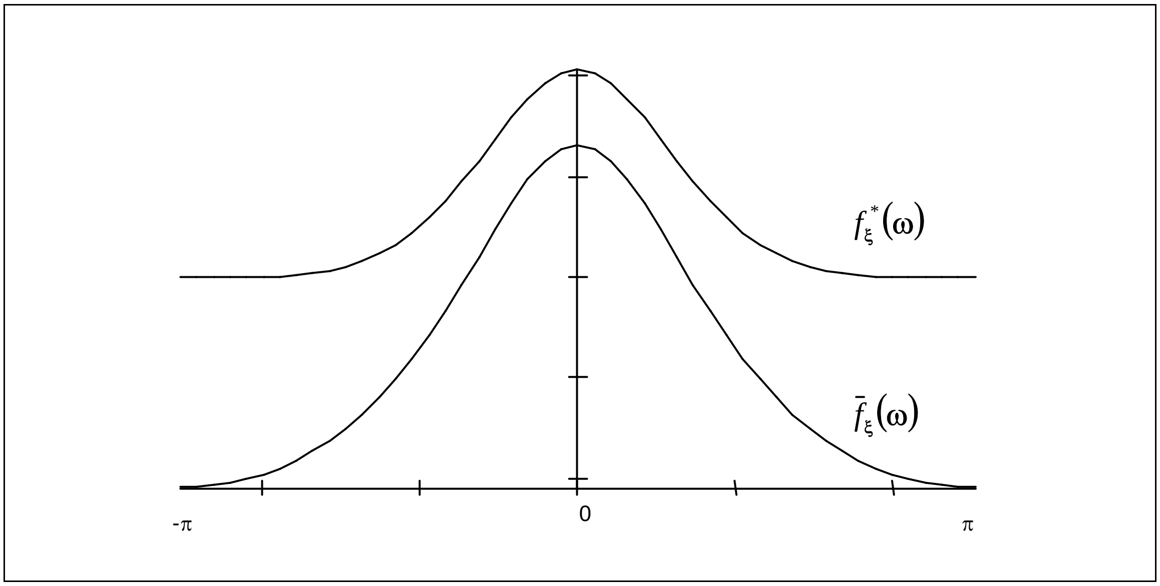

**Figure 1**

**Baseline and least favorable spectral densities**

**when $\xi_t = 0.25\xi_{t-1}+\nu_t$ , $\sigma_\nu^2 = 1$ and $\varepsilon = .01$**

## Appendix. Bode Integral Constraint

In this appendix, we provide a description of the integral constraint formula to indicate how the constant $K$ in eq. (44) is determined. To do this, one starts with $L(\omega)$, which is the transfer function that describes how shocks "move around the loop" frequency by frequency and is defined as

$$L(\omega) = H(\omega) G(\omega). \tag{86}$$

By (36) and (38), (86) may be written as

$$L(\omega) = F\left(e^{-i\omega}\right)\left(e^{i\omega} - A\left(e^{-i\omega}\right)\right)^{-1} B\left(e^{-i\omega}\right) \tag{87}$$

and so may be factored as

$$L(\omega) = c\frac{\prod_{i=1}^{m}\left(e^{-i\omega} - z_i\right)}{\prod_{i=1}^{n}\left(e^{-i\omega} - p_i\right)}. \tag{88}$$

The constant $c$ is determined by the requirement that the system is stable. The quantities $z_i$ are known as zeroes and the quantities $p_i$ are known as poles. The difference between the number of zeroes and the number of poles, $\nu = n - m$, is known as the relative degree of $L(\omega)$. We assume that $\nu \geq 1$, which will generally be the case since the policymaker cannot use current $x_t$ in choosing the control.

The following result is due to Wu and Jonckheere (1992).

**Theorem. Discrete Generalized Bode Constraint**

Assume that the controlled system is globally asymptotically stable and that $\nu \geq 1$. Then,

$$\int_{-\pi}^{\pi} \ln\left(\left|S(\omega)\right|^2\right) d\omega = 4\pi \sum_i \ln\left(\left|p_{u_i}\right|\right) \tag{89}$$

where $p_{u_i}$ denotes an unstable pole of $L(\omega)$.

It is important to note that it is possible for $x_t^{NC}$ to be either difference stationary (i.e. contain a unit root) or be explosive. What matters in our formulation is that the choice of $F(L)$ eliminates either form of nonstationarity. Lemma 5 of Wu and Jonckheere (1992) indicates that when $x_t^{NC}$ contains a unit root, (which means that $|r| = 1$, the value of the right hand side of (89) is 0 whereas if $x_t^{NC}$, it is explosive.

## Bibliography

Ash, R. and Gardner, M., 1975. Topics in Stochastic Processes. Academic Press, New York.

Brock, W. and Durlauf, S., 2004. Elements of a theory of design limits to optimal policy. The Manchester School 72, Supplement 2, 1-18.

Brock, W., Durlauf, S. and West, K., 2003. Policy evaluation in uncertain economic environments (with discussion). Brookings Papers on Economic Activity 1, 235-322.

Brock, W., Durlauf, S. and West, K., 2004. Model uncertainty and policy evaluation: some theory and empirics. Mimeo, University of Wisconsin.

Friedman, M., 1948. A monetary and fiscal framework for economic stability. American Economic Review 38, 245-264.

Giannoni, M., 2002. Does model uncertainty justify caution? Robust optimal monetary policy in a forward-looking model. Macroeconomic Dynamics 6, 111-144.

Gradshteyn, I. and Ryzhik, I., 2000. Table of Integrals, Series, and Products, Sixth Edition. Academic Press, San Diego.

Hansen, L. and Sargent, T., 2001. Acknowledging misspecification in macroeconomic theory. Review of Economic Dynamics 4, 519-35.

Hansen, L. and Sargent, T., 2003a. Robust Control and Economic Model Uncertainty. Book manuscript, Hoover Institution, Stanford University and forthcoming, Princeton University Press.

Hansen, L. and Sargent, T., 2003b. Robust control of forward-looking models. Journal of Monetary Economics 50, 581-604.

Hofbauer, J. and Sorin, S., 2002. Best response dynamics for continuous zero-sum games. Mimeo, Institut fur Mathematik, Universitat Wein, Vienna and Laboratoire d’Econometrie, Ecole Polytechnique, Paris.

Kamien, M. and Schwartz, N., 1991. Dynamic Optimization: The Calculus of Variations and Optimal Control in Economics and Management. North Holland, Amsterdam.

Kasa, K., 2000. Model uncertainty, robust policies, and the value of commitment. Macroeconomic Dynamics 6, 1, 145-166.

Kwakernaak, H. and Sivan, R., 1972. Linear Optimal Control Systems. John Wiley and Sons, New York.

Manski, C., 2004. Statistical treatment rules for heterogeneous populations. Econometrica 72, 4, 1221-1246.

Marcellino, M. and Salmon, M., 2002. Robust decision theory and the Lucas critique. Macroeconomic Dynamics 6, 1, 167-185.

Marimon, R., 1989. Stochastic turnpike theory and stationary equilibrium. Journal of Economic Theory 47, 282-306.

McKenzie, L, 1986. Optimal turnpike theorems and comparative dynamics, in: Arrow, K., Intriligator, M. (Eds.), Handbook of Mathematical Economics, Vol. 3. North-Holland, Amsterdam.

Onatski, A. and Stock, J., 2002. Robust monetary policy under model uncertainty in a small model of the U.S. economy. Macroeconomic Dynamics 6, 85-110.

Onatski, A. and Williams, N., 2003. Modeling model uncertainty. Journal of the European Economic Association 1, 1078-1122.

Sargent, T., 1987. Macroeconomic Theory, second edition. Academic Press, San Diego.

Sargent, T., 1999. Comment, in: Taylor, J. (Ed.), Monetary Policy Rules. University of Chicago Press, Chicago.

Sims, C., 2001. Pitfalls of a minimax approach to model uncertainty. American Economic Review 91, 51-4.

Svensson, L., 1996. Commentary: how should monetary policy respond to shocks while maintaining long run price stability? In: Achieving Price Stability. Federal Reserve Bank of Kansas City, pp. 181-195.

Svensson, L., 2000. Robust control made simple. Mimeo, Department of Economics, Princeton University.

Tetlow, R. and von zur Muehlen, P., 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.

Wald, A., 1950. Statistical Decision Functions. John Wiley and Sons, New York.

Woodford, M., 2003a. Interest and Prices. Princeton University Press, Princeton.

Woodford, M. 2003b. Optimal interest-rate smoothing. Review of Economic Studies 70:861-886.

Wu, B.-F. and Jonckheere, E., 1992. A simplified approach to Bode's Theorem for continuous and discrete time systems. IEEE Transactions on Automatic Control 37, 100, 1797-1802.