NBER WORKING PAPER SERIES

ROBUST OPTIMAL POLICY IN A FORWARD-LOOKING MODEL WITH PARAMETER AND SHOCK UNCERTAINTY

Marc. P. Giannoni

Working Paper 11942
http://www.nber.org/papers/w11942

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2006

This paper is based on the third chapter of my Ph.D. dissertation at Princeton University. I wish to thank especially Michael Woodford for continuous advice. I would like to thank also Jeffery Amato, Ben Bernanke, Steven Durlauf (Editor), Bart Hobijn, Thomas Laubach, Eduardo Loyo, Ernst Schaumburg, Christopher Sims, Lars Svensson, Andrea Tambalotti, Alex Wolman and an anonymous referee for valuable comments. Any remaining errors are my own. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

©2006 by Marc P. Giannoni. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Robust Optimal Policy in a Forward-Looking Model with Parameter and Shock Uncertainty
Marc P. Giannoni
NBER Working Paper No. 11942
January 2006
JEL No. C61, D81, E42, E52

ABSTRACT

This paper characterizes a robust optimal policy rule in a simple forward-looking model, when the policymaker faces uncertainty about model parameters and shock processes. We show that the robust optimal policy rule is likely to involve a stronger response of the interest rate to fluctuations in inflation and the output gap than is the case in the absence of uncertainty. Thus parameter uncertainty alone does not necessarily justify a small response of monetary policy to perturbations. However uncertainty may amplify the degree of "super-inertia" required by optimal monetary policy. We finally discuss the sensitivity of the results to alternative assumptions.

Marc P. Giannoni
Columbia Business School
3022 Broadway, Uris Hall 824
New York, NY 10027
and NBER
mg2190@columbia.edu

1 Introduction

During the last decade, economists have given increasing attention to the study of interest-rate feedback rules for the conduct of monetary policy. While some have focused on the estimation of central banks reaction functions, and the description of actual monetary policy (see, e.g., Taylor, 1993, Judd and Rudebusch, 1998, Clarida *et al.,* 2000), others have characterized optimal policy rules in the context of particular models of the economy (see, e.g., contributions collected in Taylor 1999a, Giannoni and Woodford, 2002, 2003, among many others). In reality, however, policy decisions need to be made despite considerable uncertainty about the actual functioning of the economy. Policymakers typically set their instrument without knowing the true model of the economy, and they generally do not know precisely how their policy actions will affect the variables that they care about. The prevalence of uncertainty has recently induced researchers to explore various ways to characterize desirable policy rules in the face of uncertainty (see Walsh, 2003).

A popular idea due to Brainard (1967), and emphasized by Blinder (1998) and others, is that policymakers should be cautious in the presence of uncertainty about the true parameters of a model. By "cautious" it is often meant that the instrument of monetary policy should be moved by less than in the absence of parameter uncertainty.[1] Some authors have therefore suggested that optimal policy rules that take proper account of the uncertainty surrounding model parameters should be less aggressive, and thus closer to estimated policy rules.[2] However, a number of recent studies have challenged this conventional wisdom. For instance, in Giannoni (2002), we argue that the opposite result is likely to be obtained in a simple forward-looking model that has been used in many recent studies of monetary policy. In that paper, we show that simple Taylor rules that are robust to uncertainty about structural parameters of the model may be *more responsive* to fluctuations in inflation and the output gap than the optimal Taylor in the absence of parameter uncertainty.

We call *robust optimal policy rules* policy rules that perform best in the worst-case parameter configuration, within a specified set of parameter configurations. Policy rules of this kind have

[1] As Brainard (1967) pointed out, this result holds in his setup provided that the exogenous disturbances and the parameters that relate the policy instrument to the target variable are not too strongly correlated.

[2] See Clarida *et al.* (1999), Estrella and Mishkin (1999), Hall *et al.* (1999), Martin and Salmon (1999), Svensson (1999), Rudebusch (2000), Sack (2000), Söderström (2000), and Wieland (1998), among others.

recently been advocated by Sargent (1999), Hansen and Sargent (2000, 2003, 2005), Stock (1999), Onatski and Stock (2002), and Tetlow and von zur Muehlen (2001).[3] Robust rules are designed to avoid an especially poor performance of monetary policy in the event of an unfortunate parameter configuration. They guarantee to yield an acceptable performance of monetary policy in the specified range of models.

This paper generalizes the results obtained in Giannoni (2002) in several important ways. First, instead of restricting ourselves to Taylor rules, we determine a robust optimal monetary policy rule in a family of rules that is flexible enough to implement the optimal plan, if the parameters are known with certainty. Second, we allow the model to be affected by a variety of exogenous shocks, instead of assuming a single composite exogenous perturbation. We emphasize in particular the distinction between efficient and inefficient supply shocks, as they have different welfare implications, and consider uncertainty about the relative importance of each kind of shock. Thirdly, we consider robustness of monetary policy not only to uncertainty about critical structural parameters, but also to uncertainty about the degree of persistence in the shock processes. Moreover, we emphasize the importance of deriving the model from microeconomic foundations in order to determine precisely how the exogenous disturbances are transmitted through the economy. This turns out to be important for the determination of the worst-case parameter configuration.

While it is often believed that monetary policy should be less responsive in the presence of uncertainty, we show that the opposite is likely to be true in the model considered. For a reasonable calibration of the model, the robust optimal policy rule requires the interest rate to respond more strongly to fluctuations in inflation, in changes in the output gap, and to lagged interest rates, than in the absence of uncertainty. This result depends however on the way the exogenous shocks affect the economy, and on the degree of uncertainty about the types of supply shocks.

The rest of the paper is organized as follows. The next subsection reviews briefly some of the recent literature on robust monetary policy. Section 2 describes the method used to derive the robust optimal policy rule. Section 3 presents a simple optimizing monetary model. While the model is similar to models presented in a number of recent studies, we briefly expose the microeconomic foundations of this model to specify precisely how exogenous disturbances affect

[3] Von zur Muehlen (1982) is an early study of such monetary policy rules.

the endogenous variables, when there is uncertainty about the structural parameters of the model. Section 4 characterizes both the optimal policy rule in the absence of uncertainty, and the robust optimal policy rule when there is uncertainty, and discusses the sensitivity of the results to various assumptions. Finally, section 5 concludes.

1.1 Related literature

The uncertainty faced by policymakers takes many different forms. The data that measures important economic concepts is often imperfect as it may contain measurement errors or be available only after policy decisions are made, and some of the key macroeconomic variables such as the output gap and shocks are generally not directly observed by the central bank. Optimal policy in such environments is analyzed by Aoki (2003), Orphanides (2003), and Svensson and Woodford (2003, 2004). Others, including this paper, assume that the state of the economy is perfectly observed once the shocks are realized, but that policymakers don't know the true model of the economy, so that they only have an imperfect knowledge of effect of policy actions on key economic variables. Therefore, they seek to determine policy rules that are robust to uncertainty about the correct model of the economy.

One approach, first advocated by McCallum (1988, 1999), and followed by Christiano and Gust (1999), Taylor (1999b), Levin, Wieland and Williams (1999, 2003), Levin and Williams (2003), determines policy rules that perform well across a range of models, by simulating given rules in a number of different models. Brock, Durlauf and West (2003, 2004) have made further advances by proposing a formal framework — Bayesian model averaging — and statistics grounded in decision theory to systematically evaluate alternative policy rules in the face of model uncertainty. An advantage of this approach is that it allows an analysis of model uncertainty when the models considered are potentially very different from each others. While extremely useful for understanding the effects of particular rules in various models, existing applications of this approach do actually not determine an "optimal" rule in the face of model uncertainty.

Other studies have sought to characterize optimal policy in particular classes of models, taking into account uncertainty about various aspects of the model. Researchers have for instance considered uncertainty about the parameters of the model and have used Bayesian methods to determine

the policy that minimizes the expected loss, given a prior distribution on the parameters. This approach, initially started by Brainard (1967) and developed by Chow (1975) has more recently been followed by Clarida *et al.* (1999), Wieland (1998), Estrella and Mishkin (1999), Hall *et al.* (1999), Martin and Salmon (1999), Svensson (1999), Sack (2000), Rudebusch (2001), Söderström (2000, 2002), and Kurozumi (2003) among others. Most of these studies focus on backward-looking models, and support Brainard's popular result that optimal policy should be less aggressive in the face of parameter uncertainty.[4]

Another branch of the literature has looked for robust rules that minimize a loss criterion in some worst-case scenario, within a specified set of possible scenarios. One justification for this approach is the view that uncertainty about the true model of the economy takes the form of uncertainty in the sense of Knight (1921), i.e., a situation in which the probabilities on the alternative models are not known, so that Bayesian methods cannot be used to compute the expected loss over different models.[5] Furthermore, it has been shown by Gilboa and Schmeidler (1989) that if the policymaker has multiple priors on the set of alternative models, and his preferences satisfy uncertainty aversion in addition to the axioms of standard expected utility theory, the policymaker faces a min-max problem: to minimize his loss in the worst-case scenario, i.e., when the prior distribution is the worst distribution in the set of possible distributions. Several authors have recently applied robust control theory, to derive robust monetary policies of this kind. These authors have however focused on different types of uncertainty.

For instance, Sargent (1999), Hansen and Sargent (2005), and Kasa (2002) consider very unstructured uncertainty by appending to their equations shock terms that represent model misspecifications — i.e., deviations of the model actually used from the true model — and limit uncertainty by imposing a penalty on the statistical distance (the relative entropy) between the model used and the perturbed model. They compute robust policies by minimizing a given loss criterion in the worst-case realization of the shock process that represent misspecifications. In contrast, Stock (1999), Onatski and Stock (2002), Onatski (2000a, 2000b), Onatski and Williams (2003), and

[4] One notable exception is Söderström (2002) who shows that uncertainty about the persistence of inflation induces the policymaker to respond more aggressively to shocks. Clarida *et al.* (1999) and Kurozumi (2005) consider a forward-looking model.

[5] Knight (1921) first made the distinction between "known risk," i.e. a situation in which a distribution of outcomes is known, and "uncertainty", i.e., a situation in which no known probability distribution exists.

Tetlow and von zur Muehlen (2001) consider more structured non-parametric uncertainty. They construct a non-parametric set of models around some reference model that approximates the true model of the economy, but they impose some structure on the set of possible models. They then seek to determine rules that minimize the loss for the worst possible model. These authors measure the robustness of given policy rules with the maximal size of the uncertainty set that does not include models with an indeterminate equilibrium or unstable models. While they can measure the degree of robustness of given rules, they are able to characterize the actual min-max rules only for simple types of uncertainty.

In this paper, as in Giannoni (2002), we consider uncertainty about the parameters of a structural forward-looking model. In contrast to Hansen and Sargent (2005), we maintain the rational expectations framework, by assuming that the private sector knows the true model of the economy, while the policymaker faces model uncertainty. We find the parametric treatment more intuitive, transparent, than a non-parametric approach, and believe that it allows modelers to quantify their degree of confidence more easily. This approach allows us furthermore to characterize analytically the robust rule, in sufficiently simple models. While we allow for uncertainty about a relatively small number of parameters, in the analysis below, one can in principle specify a large class of model uncertainty with parameter uncertainty, to the extent that the models can be nested parametrically. For instance, uncertainty about the variables entering particular equations, the numbers of lags of such variables, the importance of backward versus forward-looking behavior may be analyzed with parameter uncertainty. In practice, however, this approach may be more restrictive than the unstructured approach, as it may not be feasible to analyze the effects of a very large number of uncertain parameters. Approaches based on unstructured uncertainty may thus provide more convenient methods in cases in which one is worried about a wide range of possible misspecifications around a reference model. Parameter uncertainty is also not well suited to address uncertainty about models that are "disjoint" or very different from each other, and Bayesian model averaging methods advocated by Brock *et al.* (2003, 2004) may be more suited in such contexts.

So far, there is no clear answer to whether robust policy rules in the presence of uncertainty should in general be more or less aggressive than optimal rules absent model uncertainty, even

among the papers that use min-max objective functions.[6] Sargent (1999), Stock (1999), and Onatski and Stock (2002) find that robust policy requires in most cases stronger policy responses, in backward-looking models. We obtain similar results for a simple forward-looking model and in the face of parameter uncertainty. Onatski (2000b) finds robust rules to be more responsive to the output gap and less responsive to inflation in a model that involves both forward- and backward-looking elements. Whether robust policy rules should in general be more or less aggressive than optimal rules absent model uncertainty depends critically both on the model and the type of uncertainty — i.e. structured or non-structured — considered. In section 4.4, we discuss how changing various assumptions about the model can affect the results.

2 Uncertainty and Robust Optimal Monetary Policy

In reality, central banks and researchers do generally not know with certainty the true parameters of their model, in addition to not knowing the exogenous disturbances. In this paper, we assume that the parameters of the economic model are unknown to the policymaker, but remain constant over time. The policymaker commits credibly at the beginning of period 0 to a policy rule for the entire future. He chooses a policy rule to minimize some loss criterion L_0, while facing uncertainty about the true parameters of the economy. We denote by ψ the vector of coefficients that completely characterizes the policy rule, and we simply call ψ a "policy rule". We assume furthermore that the policy rules ψ are drawn from some finite-dimensional linear space $\tilde{\Psi} \in \mathbb{R}^n$.

In contrast, agents in the private sector are assumed to know the true parameters of the economy. They act optimally, i.e., in a way to maximize their utility subject to their constraints, in every period, and in every state. Specifically, we assume that the private sector may be one of many different types. Its type is determined once and for all, before period 0, and is characterized by the finite-dimensional vector of structural parameters $\theta = [\theta_1, \theta_2, ..., \theta_m]'$ defined on the compact set $\Theta \subseteq \mathbb{R}^m$. Agents in the private sector know the true type θ, but the central bank does not.

We write q_t for the vector of endogenous variables at date t, and q for the stochastic process

[6] It is sometimes believed that the results differ importantly for a Bayesian or a min-max approach. This is however not generally true. Onatski (2000a), for instance, shows that the results obtained with the min-max approach are very similar to those obtained with the Bayesian approach in the Brainard (1967) setting.

$\{q_t\}_{t=0}^{\infty}$, specifying q_t at each date as a function of the history of exogenous shocks until that date. The behavior of the private sector is determined by a set of equations for each date t, and each state. These may be written compactly as

$$\mathbf{S}(q, \theta) = \mathbf{0}. \tag{1}$$

The restrictions imposed by the commitment to the policy rule at each date can in turn be written as

$$\mathbf{P}(q, \psi) = \mathbf{0}. \tag{2}$$

A *rational expectations equilibrium* is then defined as a stochastic process $q(\psi, \theta)$ satisfying the structural equations (1) and the policy rule (2), at each date, and in every state. We restrict our attention to a subset $\Psi \subseteq \tilde{\Psi}$ of policy rules that result in a unique bounded rational expectations equilibrium, and let $q(\psi, \theta)$ denote this equilibrium.

When the structural parameters are known with certainty, the optimal monetary policy rule that is optimal relative to the subset of rules Ψ can be defined as follows.

Definition 1 *In the case of known structural parameters θ, let Ψ be a set of policy rules such that there is a unique bounded equilibrium. Then an* ***optimal monetary policy rule*** *is a vector ψ^0 that solves*

$$\min_{\psi \in \Psi} \mathrm{E}\left[L_0\left(q(\psi, \theta)\right)\right]$$

where $L_0(q)$ is the policymaker's loss function, and the unconditional expectation is taken over all possible histories of the disturbances.

To characterize parameter uncertainty, we assume that the vector θ of structural parameters lies in a given (known) compact set Θ, and that the distribution of θ is unknown. As argued in the previous section, it results from Gilboa and Schmeidler (1989) that if the policymaker has multiple priors on Θ (including the priors that any element $\theta \in \Theta$ holds with certainty), and his preferences satisfy uncertainty aversion in addition to the axioms of standard expected utility theory, the policymaker's problem is to minimize his loss in the worst-case parameter configuration. The optimal policy rule is then the robust rule defined as following.

Definition 2 *Let Ψ be a set of policy rules such that there is a unique bounded equilibrium process $q(\psi, \theta)$ for all $\psi \in \Psi, \theta \in \Theta$. In the case of parameter uncertainty, a* ***robust optimal monetary policy rule*** *is a vector ψ^* that solves*

$$\min_{\psi \in \Psi} \left\{ \max_{\theta \in \Theta} \mathrm{E}\left[L_0\left(q\left(\psi, \theta \right) \right) \right] \right\} \tag{3}$$

where $L_0(q)$ is the policymaker's loss function, and where the unconditional expectation is taken over all possible histories of the disturbances.

Given that the unknown parameter vector is in Θ, the policymaker can guarantee that the loss is no higher than the one obtained in the following "minmax" equilibrium.

Definition 3 *A* ***minmax equilibrium*** *is a bounded rational expectations equilibrium $q^* = q(\psi^*, \theta^*)$, where $\psi^* \in \Psi$ is a robust optimal monetary policy rule and θ^* maximizes the loss $\mathrm{E}\left[L_0\left(q\left(\psi^*, \theta \right) \right) \right]$ on the constraint set Θ.*

However, the equilibrium that actually realizes (given the exogenous processes) depends upon the true value of θ, and is hence unknown to the policymaker.

To characterize the robust optimal policy rule, we apply the method proposed in Giannoni (2002).[7] This method relates the solution to the problem (3) to a pure strategy Nash equilibrium (NE) of a zero-sum two-player game between a policymaker and a malevolent Nature. In this game, the policymaker chooses the policy rule $\psi^* \in \Psi$ to maximize his loss $L(\psi, \theta) \equiv \mathrm{E}\left[L_0\left(q\left(\psi, \theta \right) \right) \right]$ knowing that a malevolent Nature tries to hurt him as much as possible. Symmetrically, Nature chooses the parameter vector $\theta^* \in \Theta$ to maximize the policymaker's loss, knowing that the policymaker is going to minimize it. A NE of this game, (ψ^*, θ^*), involves a best response on the part both players. Moreover, since this is a zero-sum game, the equilibrium action of each player is a minmaximizer so that the equilibrium strategy ψ^* is a solution to (3) (see Giannoni, 2002, for additional details).

[7] Brock, Durlauf, and West (2003) propose a related approach to derive robust policy rules in the case of local uncertainty about the parameter vector. The approach adopted here, however, can be applied to situations in which the uncertainty is large.

The solution procedure involves the four following steps to characterize the robust optimal rule ψ^*.[8]

1. Optimal equilibrium for any given parameter vector θ. We determine the equilibrium process $q^*(\theta)$ that minimizes the loss $\hat{L}(q) \equiv \mathrm{E}[L_0(q)]$ subject to the restrictions imposed by the structural equations (1) for any $\theta \in \Theta$.

2. Candidate minmax equilibrium. Using $q^*(\theta)$ from step 1, we determine numerically the *candidate* worst parameter vector θ^* in the allowed set, i.e., the parameter vector that maximizes $\hat{L}(q^*(\theta))$ in the set Θ. The process $q^*(\theta^*)$ is the *candidate* minmax equilibrium.

3. Optimal policy rule. We look for a policy rule ψ^* that implements the candidate minmax equilibrium, i.e., that solves $\mathbf{P}(q^*(\theta^*), \psi^*) = \mathbf{0}$. We then verify that the policy rule ψ^* is in Ψ, i.e., that it results in a *unique* bounded equilibrium process $q(\psi^*, \theta)$ for all $\theta \in \Theta$.

4. Check for existence of *global* NE. We verify that (ψ^*, θ^*) is a *global* NE, hence that $q(\psi^*, \theta^*)$ is indeed a minmax equilibrium, by checking that the solution candidate θ^* maximizes the loss $L(\psi^*, \theta)$ on the constraint set Θ, i.e., that there is no vector $\theta^\dagger \in \Theta$ satisfying

$$L\left(\psi^*, \theta^\dagger\right) > L(\psi^*, \theta^*) \tag{4}$$

   given the policy rule ψ^*.

Steps 1 and 3 determine the policymaker's best response $\psi^* = \psi^*(\theta^*)$ to a given parameter vector θ^*. Step 2 and 4 insure in turn that θ^* is Nature's best response to ψ^*. It follows that a profile (ψ^*, θ^*) that satisfies steps 1 to 4 is a NE, and hence that ψ^* is the robust optimal rule that we are looking for. Step 4 is required to insure that the candidate worst parameter vector computed in step 2 is indeed Nature's best response to the robust optimal rule ψ^* on the whole

[8] The method presented in Giannoni (2002) is more general than the one summarized here for two reasons. First, it considers a loss function of the form $L_0(q, \theta)$, where the second argument allows the coefficients of the loss function to be functions of the parameter vector θ. Second, it allows one to characterize robust optimal rules in restricted families of policy rules. As these restricted families of rules impose restrictions besides (1) and (2) on the space of possible processes, the space of possible processes is parametrized by an alternative parameter vector f. We don't need to consider this complication here, as the family of policy rules that we consider below does not impose any additional restrictions besides (1) and (2).

constraint set Θ, so that (ψ^*, θ^*) is not only a *local* NE — i.e., a situation in which each player's strategy is at least locally a best response to the other player's strategy — but also a *global* NE. Note that a global NE may not exist, even though a robust optimal rule should still exist. However, in applications such as the one in section 4, a global NE will exist.

While steps 2 and 4 require a numerical maximization of the loss function with respect to θ, on the set Θ, it is simpler to characterize the robust optimal rule following the four steps mentioned here, than trying to solve (3) directly. Indeed, solving (3) would require maximizing the loss function over θ for any given policy rule ψ, until the robust rule ψ^* is obtained. In addition, the solution procedure proposed here may allow one to obtain an analytical characterization of the robust rule as will be the case in section 4.[9]

3 A Simple Optimizing Model for Monetary Policy Analysis

The model that characterizes the behavior of the private sector is a variant of the "new Keynesian" or "new synthesis" model presented, e.g., in Clarida *et al.* (1999) and Woodford (2003). In order to understand precisely how the shocks affect the economy, we briefly describe the model that characterizes the private sector's behavior, and then turn to the objective of monetary policy.

3.1 Underlying Structural Model

We assume that there exists a continuum of households indexed by j and distributed uniformly on the $[0,1]$ interval. Each household j consumes all of the goods and supplies a single differentiated good. It seeks to maximize its lifetime expected utility given by

$$\mathrm{E}_0 \left\{ \sum_{t=0}^{\infty} \beta^t \left[u\left(C_t^j; \xi_t\right) + \chi\left(M_t^j / P_t; \xi_t\right) - v\left(y_t\left(j\right); \xi_t\right) \right] \right\} \tag{5}$$

where $\beta \in (0,1)$ is the household's discount factor (assumed to be equal for each household), M_t^j is the amount of money balances held at the end of period t, $y_t\left(j\right)$ is the household's supply of its

[9] Note that if we compute the worst vector θ^* by maximizing directly $L\left(q\left(\psi^*\left(\theta\right), \theta\right)\right)$ with respect to $\theta \in \Theta$, we would obtain the solution to $\max_{\theta \in \Theta} \left\{ \min_{\psi \in \Psi} \mathrm{E}\left[L_0\left(q\left(\psi, \theta\right)\right)\right] \right\}$, and not necessarily the parameter vector θ that solves (3). The solution to both problems is however the same provided that it is part of a global NE. Our four-step procedure guarantees that we obtain the robust policy rule that we are looking for, provided that a global NE exists.

good, C_t^j is an index of the household's consumption of each of the differentiated goods defined by

$$C_t^j \equiv \left[\int_0^1 c_t^j(z)^{\frac{\varphi_t-1}{\varphi_t}} dz\right]^{\frac{\varphi_t}{\varphi_t-1}}, \tag{6}$$

and P_t is the corresponding price index. The consumption index aggregates consumption of each good, $c_t^j(z)$, with an elasticity of substitution between goods, $\varphi_t > 1$, at each date. In contrast to Dixit and Stiglitz (1977) however, we let the elasticity of substitution vary exogenously over time. As will appear more clearly below, such perturbations to the elasticity of substitution imply time variation in the price elasticity of demand of each good, and variations of the desired markup. The stationary vector ξ_t represents disturbances to preferences. For each value of ξ, the functions $u(\cdot;\xi)$ and $\chi(\cdot;\xi)$ are assumed to be increasing and concave, while the disutility from supplying goods, $v(\cdot;\xi)$, is increasing and convex.

Expenditure minimization and market clearing imply that the demand for each good j is given by

$$y_t(j) = Y_t\left(\frac{p_t(j)}{P_t}\right)^{-\varphi_t} \tag{7}$$

where $p_t(j)$ is the price of good j, and $Y_t = C_t \equiv \int_0^1 C_t^j dj$ represents aggregate demand at date t.

We assume that financial markets are complete so that risks are efficiently shared. It follows that all households face an identical intertemporal budget constraint, and choose identical state-contingent plans for consumption, and money balances. We may therefore drop the index j on those variables.

Each household maximizes (5) subject to its budget constraint, and the constraint that it satisfies the demand for its good (7). It follows that the optimal intertemporal allocation of consumption satisfies a familiar Euler equation of the form

$$\frac{1}{1+i_t} = \mathrm{E}_t\left\{\frac{\beta u_c\left(Y_{t+1};\xi_{t+1}\right)}{u_c\left(Y_t;\xi_t\right)}\frac{P_t}{P_{t+1}}\right\}, \tag{8}$$

where i_t denotes the nominal interest rate on a riskless one-period nominal bond purchased in period t. We will consider a log-linear approximation of this relationship about the steady state where the exogenous disturbances take the values $\xi_t = 0$ and where there is no inflation. We

let $\bar{Y}$ and $\bar{\imath}$ be the constant values of output and nominal interest rate in that steady state, and define the percent deviations $\hat{Y}_t \equiv \log\left(Y_t/\bar{Y}\right)$, $\hat{\imath}_t \equiv \log\left(\frac{1+i_t}{1+\bar{\imath}}\right)$, $\pi_t \equiv \log\left(P_t/P_{t-1}\right)$. The log-linear approximation to (8) is

$$\hat{Y}_t = \mathrm{E}_t\hat{Y}_{t+1} - \sigma^{-1}\left(\hat{\imath}_t - \mathrm{E}_t\pi_{t+1}\right) + \sigma^{-1}\delta_t \tag{9}$$

where $\sigma \equiv -\frac{u_{cc}\bar{C}}{u_c} > 0$ represents the inverse of the intertemporal elasticity of substitution in private expenditures, and where

$$\delta_t \equiv \frac{u_{c\xi}}{u_c}\left(\xi_t - \mathrm{E}_t\xi_{t+1}\right) \tag{10}$$

represents exogenous disturbances to (9). Equation (9), which represents the demand side of the economy, is often called the "intertemporal IS equation" as it relates negatively desired expenditures to the real interest rate. We assume that δ_t is independent of σ.[10]

Monetary policy has real effects in this model because prices do not respond immediately to perturbations. Specifically, we assume as in Calvo (1983) that only a fraction $1-\alpha$ of suppliers may change their prices at the end of any given period, regardless of the the time elapsed since the last change. Because of monopolistic competition, each household chooses the optimal prices $\{p_t(j)\}$, taking as given the evolution of aggregate demand and the price level, that determine the location of the demand for its product (7). Each supplier that changes its price in period t chooses its new price to maximize the present discounted value of its expected future profits. Log-linearizing the resulting first-order conditions, we obtain the following aggregate supply equation

$$\pi_t = \kappa\left(\hat{Y}_t - \hat{Y}_t^n\right) + \beta\mathrm{E}_t\pi_{t+1}, \tag{11}$$

where $\kappa > 0$, and $\hat{Y}_t^n$ represents the *natural* rate of output, i.e., the percentage deviations from steady-state of the level of output that would obtain with perfectly flexible prices (see, e.g., Woodford, 2003, for details). As further shown in the appendix of Giannoni (2000), the natural rate of output satisfies

$$\hat{Y}_t^n = \frac{1}{\omega+\sigma}\left(\frac{u_{c\xi}}{u_c}\xi_t - \frac{v_{y\xi}}{v_y}\xi_t - \mu_t\right), \tag{12}$$

[10] This is true, for instance, for any utility function of the form $u(C,\xi) = v(C)\cdot w(\xi)$ where v and w are independent of each other, since $\sigma = -\frac{v''\bar{C}}{v'}$, and $\delta_t = \frac{w'}{w}\left(\xi_t - \mathrm{E}_t\xi_{t+1}\right)$ in this case.

where $\omega > 0$ represents the elasticity of each firm's real marginal cost with respect to its own supply and μ_t represents percent deviations of the desired markup $\varphi_t/(\varphi_t - 1)$ from steady state. Note that while the natural rate of output depends upon both supply and demand exogenous real perturbations, it is completely independent of monetary policy. Because of market power, however, steady-state level of output is inefficiently low. As the percent deviations of the *efficient* rate of output — i.e., the equilibrium rate of output that would obtain in the absence of price rigidities and market power — are given by $\hat{Y}_t^e = \frac{1}{\omega+\sigma}\left(\frac{u_{c\xi}}{u_c}\xi_t - \frac{v_{y\xi}}{v_y}\xi_t\right)$, exogenous time variation in the desired markup results in deviations of the efficient rate of output from the natural rate given by

$$\hat{Y}_t^e - \hat{Y}_t^n = \frac{1}{\omega+\sigma}\mu_t.$$

As we will evaluate monetary policy in terms of deviations of output from its efficient level, it will be convenient to define the "output gap" as

$$x_t \equiv \hat{Y}_t - \hat{Y}_t^e. \tag{13}$$

Using this, we can rewrite the two structural equations (9) and (11) as

$$x_t = \mathrm{E}_t x_{t+1} - \sigma^{-1}\left(\hat{\imath}_t - \mathrm{E}_t\pi_{t+1}\right) + \frac{\omega}{(\omega+\sigma)\,\sigma}\delta_t + \frac{1}{\omega+\sigma}\varepsilon_t \tag{14}$$

$$\pi_t = \kappa\left(x_t + \frac{1}{\omega+\sigma}\mu_t\right) + \beta\mathrm{E}_t\pi_{t+1}, \tag{15}$$

where δ_t is the demand shock defined in (10), and where

$$\varepsilon_t \equiv \frac{v_{y\xi}}{v_y}\left(\xi_t - \mathrm{E}_t\xi_{t+1}\right)$$

is an adverse "efficient" supply shock. We suppose that the vector of shocks $u_t \equiv [\delta_t, \varepsilon_t, \mu_t]'$ satisfies $\mathrm{E}(u_t) = 0$, and that these perturbations are independent of the parameters σ, κ, or ω.[11]

As in Giannoni (2000), we call the exogenous disturbance to the aggregate supply equation, μ_t, an "*inefficient supply shock*" since it represents a perturbation to the natural rate of output that is

[11] Again, ε_t is independent of ω if, for instance, the disutility of supplying goods is of the form $v(y,\xi) = \varpi(y)\cdot\nu(\xi)$. (See footnote 10.)

not efficient. While μ_t represents fluctuations in the desired markup, this term may alternatively represent variations in distortionary tax rates, or variations in the degree of market power of workers. We prefer to call μ_t an "inefficient supply shock" rather than a "cost-push shock" as is often done in the literature (see, e.g., Clarida *et al.*, 1999), because perturbations that affect inflation by changing costs may well change the efficient rate of output as well as the natural rate of output. It follows that cost shocks are represented in our model by changes in x_t rather than μ_t.

Many recent studies have emphasized the role of the "natural" or "efficient" rate of interest for evaluating the stance of monetary policy (see, e.g., Blinder 1998, Woodford, 2003). The efficient rate of interest, i.e., the equilibrium real interest rate that would equate output to the efficient rate of output, $\hat{Y}_t^e$, is defined here as

$$r_t^e \equiv \frac{\omega}{\omega+\sigma}\delta_t + \frac{\sigma}{\omega+\sigma}\varepsilon_t. \tag{16}$$

Equation (14) can then be rewritten as

$$x_t = \mathrm{E}_t x_{t+1} - \sigma^{-1}\left(\hat{\imath}_t - \mathrm{E}_t\pi_{t+1} - r_t^e\right). \tag{17}$$

It is clear from (16) that the efficient rate of interest depends both on demand shocks δ_t and efficient supply shocks ε_t. It follows from (17) that monetary policy is expansive or restrictive only insofar as the equilibrium real interest rate is below or above the efficient rate. If the central bank was perfectly tracking the path of r_t^e, then the output gap would be zero at all times, and inflation would only depend on fluctuations in μ_t.[12]

3.2 Monetary Policy

We now turn to the objective of monetary policy. The policymaker is assumed to have the following loss function

$$L_0 = \mathrm{E}_0\left\{(1-\beta)\sum_{t=0}^{\infty}\beta^t\left[\pi_t^2 + \lambda_x\left(x_t - x^*\right)^2 + \lambda_i\hat{\imath}_t^2\right]\right\} \tag{18}$$

[12] There is an additional first-order condition that determines the optimal holdings of monetary balances as a function of equilibrium consumption (or output), the nominal interest rate, and the price level. When monetary policy determines the nominal interest rate, as is the case here, this condition can be omitted as it has no effect on the equilibrium values of inflation, output, and nominal interest rate. The presence of real balances in the utility function (5) matters however for the determination of the loss function below.

where $\lambda_x, \lambda_i > 0$, are weights placed on the stabilization of the output gap and the nominal interest rate, and where $x^* \geq 0$ represents some optimal level of the output gap. (Note that we implicitly assume that the optimal levels of both inflation and the interest rate are zero). As in many studies of monetary policy, we assume that the policymaker seeks to stabilize fluctuations in inflation and in the output gap. We furthermore assume that he also cares about the variability of the nominal interest rate, as a result of transaction frictions. Friedman (1969) has argued that high nominal interest rates involve welfare costs of transactions. Whenever the deadweight loss is a convex function of the distortion, then it is desirable to reduce not only the level but also the variability of the nominal interest rate (see Woodford, 1990, 2003). Such a loss criterion can finally obtained as a second-order Taylor approximation to the utility function of the household's lifetime utility (5) in equilibrium, when the parameters are known with certainty. We will assume that the policymaker minimizes the unconditional expectation of the above loss criterion, $\mathrm{E}[L_0]$, where the expectation is taken with respect to the stationary distribution of the shocks. As a result, optimal policy will be independent of the initial state.

We characterize monetary policy in terms of interest-rate rules. Specifically, we assume that the policymaker commits credibly at the beginning of period 0 to a feedback rule of the form

$$\hat{\imath}_t = \mathcal{P}_t\left(\pi_t, \pi_{t-1}, ..., x_t, x_{t-1}, ..., \hat{\imath}_{t-1}, \hat{\imath}_{t-2}, ..., u_t, u_{t-1}, ...\right) \tag{19}$$

for each date $t \geq 0$. The policymaker determines the functions $\mathcal{P}_t(\cdot)$, $t = 0, 1, 2, ...$ to minimize the loss $\mathrm{E}[L_0]$ subject to the structural equations (14) and (15). As the objective is quadratic and the constraints are linear in all variables, we may without loss of generality restrict our attention to linear functions $\mathcal{P}_t(\cdot)$. Using the notation of section 2, we denote by ψ the finite-dimensional vector of coefficients that completely characterizes $\{\mathcal{P}_t(\cdot)\}_{t=0}^{\infty}$, and we call ψ a "policy rule".

3.3 Calibration

The model considered here is very similar to a simplified version of the econometric model that Rotemberg and Woodford (RW) (1997, 1999) have estimated for the US economy. The structural equations in RW correspond to (9) and (11) only when conditioned upon information available

two quarters earlier.[13] We will use their estimates to calibrate our model, in the baseline case. RW calibrate β, setting it at 0.99. They estimate $\sigma = .1571$, $\kappa = .0238$. The standard errors (se) for these parameters are respectively 0.0328 and 0.0035. These numbers were computed for the RW model using the estimation method explained in Amato and Laubach (2003).[14] Finally, RW calibrate ω, setting it at 0.4729. As we will consider uncertainty also about ω we will assume that the standard error is 0.0946, corresponding to 20% of the calibrated value (which is approximately in line with the uncertainty about σ and κ). We assume that the uncertainty about the critical structural parameters is given by the approximate 95% intervals

$$\begin{aligned}
[\underline{\sigma}, \bar{\sigma}] &= [\sigma - 2se_\sigma, \sigma + 2se_\sigma] = [0.0915, 0.2227] \\
[\underline{\kappa}, \bar{\kappa}] &= [\kappa - 2se_\kappa, \kappa + 2se_\kappa] = [0.0168, 0.0308] \\
[\underline{\omega}, \bar{\omega}] &= [\omega - 2se_\omega, \omega + 2se_\omega] = [0.2837, 0.6621] .
\end{aligned}$$

In Section 4.4, we consider the case in which there is much more uncertainty about these parameters, so that these intervals are considerably wider. For simplicity, we assume that β is known with certainty. We now turn to the calibration of the variance-covariance matrix of the exogenous disturbances. RW estimate the process for the exogenous variables $\hat{G}_t$, $\hat{Y}_t^S$ in their model. This process is given by

$$\begin{aligned}
\begin{bmatrix} \hat{G}_{t+1} \\ \hat{Y}_{t+1}^S \end{bmatrix} &= \begin{bmatrix} c_1 \\ c_2 \end{bmatrix} \bar{Z}_{t-1} + \begin{bmatrix} d_1 \\ d_2 \end{bmatrix} \bar{e}_t \\
\bar{Z}_t &= B\bar{Z}_{t-1} + U\bar{e}_t
\end{aligned}$$

[13] When conditioning both the intertemporal IS equation and the aggregate supply equation in RW (1997) upon information available at $t-2$, we obtain

$$\begin{aligned}
\mathrm{E}_{t-2}\mathcal{Y}_t &= \mathrm{E}_{t-2}\mathcal{Y}_{t+1} - \sigma^{-1}\mathrm{E}_{t-2}\left(\hat{R}_t - \hat{\pi}_{t+1}\right) + \mathrm{E}_{t-2}\left(\hat{G}_t - \hat{G}_{t+1}\right) \\
\mathrm{E}_{t-2}\hat{\pi}_t &= \kappa\mathrm{E}_{t-2}\left(\mathcal{Y}_t - \hat{Y}_t^S\right) + \beta\mathrm{E}_{t-2}\hat{\pi}_{t+1},
\end{aligned}$$

where $\mathcal{Y}_t, \hat{\pi}_t, \hat{R}_t$ represent respectively output, inflation, and the nominal interest rate expressed as percentage deviations from steady state in RW (1997), and where $\hat{G}_t$ is an exogenous variable representing autonomous changes in demand, and $\hat{Y}_t^S$ represents exogenous disturbances to the aggregate supply equation. Defining $\hat{Y}_t \equiv \mathrm{E}_t\mathcal{Y}_{t+2}$, $\pi_t \equiv \mathrm{E}_t\hat{\pi}_{t+2}$, $\hat{\imath}_t \equiv \mathrm{E}_t\hat{R}_{t+2}$, $\sigma^{-1}\frac{u_{c\xi}}{u_c}\xi_t \equiv \mathrm{E}_t\hat{G}_{t+2}$, and $\hat{Y}_t^n \equiv \mathrm{E}_t\hat{Y}_{t+2}^S$, we obtain (9) and (11).

[14] I am grateful to Thomas Laubach for providing me with these numbers.

where $\mathrm{E}_t\bar{e}_{t+j} = 0$ for all $j > 0$, and the variance-covariance matrix of the state vector $\bar{Z}_t$ is Ω. The variables $\mathrm{E}_t\hat{G}_{t+2}$ and $\mathrm{E}_t\hat{Y}^S_{t+2}$ in their model correspond respectively to $\sigma^{-1}\frac{u_{c\xi}}{u_c}\xi_t$ and $\hat{Y}^n_t$ in our model. It follows that the process for δ_t is given by

$$\delta_t = \sigma\mathrm{E}_t\left(\hat{G}_{t+2} - \hat{G}_{t+3}\right) = \sigma c_1\left(\bar{Z}_t - \mathrm{E}_t\bar{Z}_{t+1}\right) = \sigma c_1\left(I - B\right)\bar{Z}_t.$$

Let us define the supply shock

$$s_t \equiv \mu_t + \frac{v_{y\xi}}{v_y}\xi_t$$

We know from (12) that $s_t = \frac{u_{c\xi}}{u_c}\xi_t - (\omega + \sigma)\hat{Y}^n_t$. It follows from the above equations that

$$\begin{aligned} s_t &= \sigma\mathrm{E}_t\hat{G}_{t+2} - (\omega + \sigma)\mathrm{E}_t\hat{Y}^S_{t+2} \\ &= h\bar{Z}_t \end{aligned}$$

where $h \equiv \sigma c_1 - (\omega + \sigma)c_2$. While we can characterize the process for s_t, we don't have enough information to determine the split between the efficient component $\frac{v_{y\xi}}{v_y}\xi_t$, and the inefficient supply shock μ_t. We therefore simply assume that $\mu_t = \nu s_t$ and $\frac{v_{y\xi}}{v_y}\xi_t = (1 - \nu)s_t$, where ν is some constant between 0 and 1. It follows that the processes for the two supply shocks are given by

$$\varepsilon_t = (1 - \nu)h\left(\bar{Z}_t - \mathrm{E}_t\bar{Z}_{t+1}\right) = (1 - \nu)h\left(I - B\right)\bar{Z}_t,$$

and

$$\mu_t = \nu h\bar{Z}_t.$$

As a result, the variance-covariance matrix of the vector of exogenous disturbances u_t is given by

$$\mathrm{E}\left(u_t u_t'\right) = \begin{bmatrix} \sigma c_1\left(I - B\right) \\ (1 - \nu)h\left(I - B\right) \\ \nu h \end{bmatrix} \Omega \begin{bmatrix} \sigma c_1\left(I - B\right) \\ (1 - \nu)h\left(I - B\right) \\ \nu h \end{bmatrix}'. \tag{20}$$

We will compute the covariance matrix for different values of ν. Below we will consider uncertainty about ν, knowing only that ν lies between 0 and 1.[15] Finally, we will assume that the three exogenous shocks follow an AR(1) process, with coefficients of serial correlation of $\rho_\delta, \rho_\varepsilon, \rho_\mu$. Woodford (1999) argues that the coefficient of autocorrelation of the natural rate of interest is 0.35. We will consider as a benchmark the case in which $\rho_\delta = \rho_\varepsilon = \rho_\mu = 0.35$, but we will also consider the case in which there is uncertainty about the coefficients of autocorrelation, allowing their values to be anywhere in the $[0, 0.8]$ interval. The benchmark parameters are summarized in Table 1.

4 Robust Optimality within a Flexible Class of Interest-Rate Rules

We now turn to the characterization of optimal monetary policy within a flexible class of interest-rate rules $\tilde{\Psi}$ that allow the instrument to respond to past variables. We define $\tilde{\Psi}$ as the set of policy rules

$$\psi = [\psi_\pi, \psi_x, \psi_{i1}, \psi_{i2}]'$$

satisfying

$$\hat{\imath}_t = \psi_\pi \pi_t + \psi_x (x_t - x_{t-1}) + \psi_{i1} \hat{\imath}_{t-1} + \psi_{i2} \hat{\imath}_{t-2} \tag{21}$$

at all dates $t \geq 0$.[16] As will become clear below, the set $\tilde{\Psi}$ is flexible enough to include a fully optimal rule in the case of any parameter vector $\theta \in \Theta$ (if the parameters were known with certainty), though it is still specific enough to contain only one rule consistent with the optimal plan in any such case. Moreover this class of rules includes recent descriptions of actual monetary policy such as the one proposed by Judd and Rudebusch (1998). We start with the characterization of the optimal plan for a given θ, and propose an interest-rate rule that implements that plan. We next determine the minmax equilibrium, and the robust optimal policy rule that implements it.

[15] The assumption that both μ_t and $\frac{v_{y\xi}}{v_y}\xi_t$ are proportional to s_t may seem unappealing as it implies that these variables are perfectly correlated, as long as $0 < \nu < 1$. However, as we will see below, once we consider uncertainty about ν, the variance-covariance matrix of u_t that matters is actually either the one for which $\nu = 0$ or the one for which $\nu = 1$.

[16] As we evaluate monetary policy regardless of specific initial conditions, the policy rule is assumed to be independent of the values the endogenous variables might have taken before it was implemented. Specifically, we assume that the policymaker considers the initial values as satisfying $i_{-2} = i_{-1} = x_{-1} = 0$, whether they actually do or not. Equivalently, we could assume that the policy rule satisfies $i_0 = \psi_\pi \pi_0 + \psi_x x_0$, $i_1 = \psi_\pi \pi_1 + \psi_x (x_1 - x_0) + \psi_{i1} i_0$, and (21) at all dates $t \geq 2$.

4.1 Optimal Plan with Given Parameters

To characterize the optimal plan for a given parameter vector $\theta \in \Theta$, we determine the stochastic process $q^*(\theta)$ of endogenous variables that minimizes the unconditional expectation of the loss criterion (18) subject to the constraints (14) and (15) at all dates $t \geq 0$, and in every state that may occur at date t, i.e., for every possible history of the disturbances until that date. In terms of the notation laid out in section 2, we determine the stochastic process $q^*(\theta)$ that minimizes the loss $\hat{L}(q)$ subject to the restrictions (1) imposed by the structural equations (14) and (15). The policymaker's Lagrangian can be written as

$$\begin{aligned}\mathcal{L} &= \mathrm{E}\Bigg\{\mathrm{E}_0(1-\beta)\sum_{t=0}^{\infty}\beta^t\Bigg(\left[\pi_t^2+\lambda_x(x_t-x^*)^2+\lambda_i\hat{\imath}_t^2\right] \\ &\quad +2\phi_{1t}\left[x_t-x_{t+1}+\sigma^{-1}(\hat{\imath}_t-\pi_{t+1})-\frac{\omega}{(\omega+\sigma)\sigma}\delta_t-\frac{1}{\omega+\sigma}\varepsilon_t\right] \\ &\quad +2\phi_{2t}\left[\pi_t-\kappa\left(x_t+\frac{1}{\omega+\sigma}\mu_t\right)-\beta\pi_{t+1}\right]\Bigg)\Bigg\}. \qquad (22)\end{aligned}$$

The first-order necessary conditions with respect to π_t, x_t, and $\hat{\imath}_t$ are

$$\pi_t-(\beta\sigma)^{-1}\phi_{1t-1}+\phi_{2t}-\phi_{2t-1} = 0 \qquad (23)$$

$$\lambda_x(x_t-x^*)+\phi_{1t}-\beta^{-1}\phi_{1t-1}-\kappa\phi_{2t} = 0 \qquad (24)$$

$$\lambda_i\hat{\imath}_t+\sigma^{-1}\phi_{1t} = 0 \qquad (25)$$

at each date $t \geq 0$, and for each possible state. In addition, we have the initial conditions

$$\phi_{1,-1}=\phi_{2,-1}=0 \qquad (26)$$

indicating that the policymaker has no previous commitment at time 0. Note that since the objective function is convex in q, and the constraints are linear in q, (14), (15), and (23) – (25) at all dates t, together with the initial condition (26) are not only necessary but also sufficient to determine the bounded optimal plan $\{\pi_t, x_t, \hat{\imath}_t, \phi_{1t}, \phi_{2t}\}$. In the steady-state, i.e., in the absence of perturbations, (14), (15), and (23) – (26) reveal that the endogenous variables remain constant at

the values

$$\begin{aligned}\pi^{op} &= x^{op} = i^{op} = \phi_1^{op} = 0 \\ \phi_2^{op} &= -\frac{\lambda_x}{\kappa} x^*.\end{aligned}$$

It will be convenient to replace ϕ_{2t} with $\hat{\phi}_{2t} \equiv \phi_{2t} - \phi_2^{op}$ so that the constant drops out of (24). Using (25) to substitute for the interest rate, we can rewrite the dynamic system (14), (15), (23), and (24) in matrix form as

$$\mathrm{E}_t \begin{bmatrix} z_{t+1} \\ \phi_t \end{bmatrix} = M \begin{bmatrix} z_t \\ \phi_{t-1} \end{bmatrix} + m u_t, \tag{27}$$

where $z_t \equiv [\pi_t, x_t]'$, $\phi_t \equiv \left[\phi_{1t}, \hat{\phi}_{2t}\right]'$, $u_t \equiv [\delta_t, \varepsilon_t, \mu_t]'$, and M and m are matrices of coefficients. Following Blanchard and Kahn (1980), this dynamic system has a unique bounded solution (given a bounded process $\{u_t\}$) if and only if the matrix M has exactly two eigenvalues outside the unit circle. Investigation of the matrix M reveals that if a bounded solution exists, it is unique.[17] In this case the solution for the endogenous variables can be expressed as

$$q_t = D\phi_{t-1} + \sum_{j=0}^{\infty} d_j \mathrm{E}_t u_{t+j} \tag{28}$$

where $q_t \equiv [\pi_t, x_t, \hat{\imath}_t]'$, and the Lagrange multipliers follow the law of motion

$$\phi_t = N\phi_{t-1} + \sum_{j=0}^{\infty} n_j \mathrm{E}_t u_{t+j} \tag{29}$$

for some matrices D, N, d_j, n_j that depend upon the parameters of the model. Woodford (1999) has emphasized that in the optimal plan with given structural parameters, the endogenous variables should depend not only upon expected future values of the disturbances, but also upon the predetermined variables ϕ_{t-1}. This dependence indicates that optimal monetary policy should involve inertia in the interest rate, regardless of the possible inertia in the exogenous perturbations. In fact, as argued by Woodford (1999), policymakers who choose optimal actions by disregarding their past

[17] The matrix M has two eigenvalues with modulus greater than $\beta^{-1/2}$ and two with modulus smaller than this.

actions and past states of the economy, don't achieve the best equilibrium when the private sector is forward-looking. The central bank should realize that the evolution of its goal variables depends not only upon its current actions, but also upon how the private sector foresees future monetary policy. It should therefore act in a way that affects the response of the private sector appropriately. As will become clearer below, it can do so by committing itself to a rule of the kind (21).

Figure 1 plots with solid lines the optimal response of the interest rate, inflation, and the output gap to an unexpected demand shock, in the baseline calibration.[18] The disturbance δ_t unexpectedly increases by 1 at date 0 and is expected to return to steady state following an AR(1) process with a coefficient of autocorrelation of $\rho_\delta = 0.35$. The path that the efficient rate of interest (16) is expected follow is indicated by the dashed-dotted line in the upper panel. Similar impulse responses would be generated by an adverse efficient supply shock, i.e., an increase in ε_t. While the policymaker could in principle completely stabilize the output gap and inflation, by tracking the path of the efficient rate of interest, it is optimal to increase the nominal interest rate by less than the efficient rate of interest at the period of the shock because the policymaker also wants to dampen fluctuations in the nominal rate of interest. As monetary policy is relatively expansionary, inflation and the output gap increase in response to the perturbation. The short-term interest is also more inertial than the efficient rate. Inertia in monetary policy is especially desirable here because it induces the private sector to expect future negative output gaps which in turn have a disinflationary effect. Therefore, by acting in an inertial way, the policymaker can offset the inflationary impact of the shock without having to raise the short-term interest much. Qualitatively similar figures would be obtained for different degrees of serial correlation in the perturbations.

Figure 2 displays with solid lines the optimal response of the endogenous variables to an unexpected inefficient supply shock μ_t. Specifically, we assume that the desired markup increases unexpectedly by one percentage point at date 0, and is expected to return to steady-state according to AR(1) process with coefficient of autocorrelation $\rho_\mu = 0.35$. Figure 2 reveals that it is optimal to slightly raise the nominal interest rate. This helps maintaining the output gap (and output since there is no change in $\hat{Y}_t^e$) below steady state for several periods. As a result, the private sector

[18] The impulse responses of all variables are reported in annual terms. Therefore, the responses of $\hat{\imath}_t$ and π_t are multiplied by 4.

expects a slight deflation in the future, which removes some inflationary pressure already at the time of the shock.

4.2 Optimal Interest-Rate Rule with Given Parameters

We now turn to the determination of an optimal interest rate rule, namely the policy rule in the family (21) that implements the optimal plan, for given structural parameters θ. We solve (25) for ϕ_{1t} and (24) for $\hat{\phi}_{2t}$, and use the resulting expressions to substitute for the Lagrange multipliers in (23). This yields

$$\hat{\imath}_t = \frac{\kappa}{\lambda_i \sigma}\pi_t + \frac{\lambda_x}{\lambda_i \sigma}\left(x_t - x_{t-1}\right) + \left(1 + \frac{\kappa}{\beta\sigma} + \beta^{-1}\right)\hat{\imath}_{t-1} - \beta^{-1}\hat{\imath}_{t-2} \tag{30}$$

for all $t \geq 0$.[19] As this equilibrium condition relates the endogenous variables in the optimal plan, the policy rule

$$\psi^*(\theta) = \left[\frac{\kappa}{\lambda_i \sigma}, \frac{\lambda_x}{\lambda_i \sigma}, \left(1 + \frac{\kappa}{\beta\sigma} + \beta^{-1}\right), -\beta^{-1}\right]' \tag{31}$$

satisfies the restrictions $\mathbf{P}\left(q^*(\theta), \psi^*(\theta)\right) = \mathbf{0}$. Furthermore, since the endogenous variables entering (30) minimize the loss criterion $\hat{L}(q)$ subject to the constraints (1) in the optimal plan, the following lemma guarantees that $\psi^*(\theta)$ is an optimal rule for any given $\theta \in \Theta$, provided that it results in a unique bounded equilibrium.

Lemma 4 *Suppose that $q^*(\theta)$ minimizes $\hat{L}(q)$ subject to (1) for any given $\theta \in \Theta$, and that there exists $\psi^*(\theta) \in \Psi$ that solves $\mathbf{P}\left(q^*(\theta), \psi^*(\theta)\right) = \mathbf{0}$ for all $\theta \in \Theta$. Then $\psi^*(\theta) \in \arg\min_{\psi \in \Psi} L(\psi, \theta)$.*

Proof. First note that since $\psi^*(\theta) \in \Psi$, the latter policy rule results in a unique bounded equilibrium. Suppose as a way of contradiction that there exists a policy rule $\psi^\dagger(\theta) \in \Psi$, $\psi^\dagger(\theta) \neq \psi^*(\theta)$, satisfying $L\left(\psi^\dagger(\theta), \theta\right) < L\left(\psi^*(\theta), \theta\right)$. By definition of $L(\cdot)$ and $\hat{L}(\cdot)$ we have $L(\psi, \theta) = \hat{L}\left(q(\psi, \theta)\right)$ for all $\psi \in \Psi, \theta \in \Theta$, so that $\hat{L}\left(q\left(\psi^\dagger(\theta), \theta\right)\right) < \hat{L}\left(q\left(\psi^*(\theta), \theta\right)\right) = \hat{L}\left(q^*(\theta)\right)$. But then $q^*(\theta)$ cannot minimize $\hat{L}(q)$ subject to (1). ■

[19] See remarks in footnote 16.

The dynamic system obtained by combining (14), (15), and (30), has the property of system (27) that, if any bounded solution exists, it is unique.[20] Moreover it can be shown, at least in the baseline parametrization, and for all values $\theta \in \Theta$ of our example, that $\psi^*(\theta)$ is the *unique* optimal policy rule in the set $\tilde{\Psi}$ (see Appendix B1 in Giannoni, 2000).

Notice that this rule makes no reference to any of the exogenous shocks. It achieves the minimal loss *regardless* of the processes that describe the evolution of δ_t, ε_t, and μ_t, provided that the latter processes are stationary (bounded).[21] If we would allow for a broader class of policy rules than $\tilde{\Psi}$, other interest-rate feedback rules may implement the same optimal plan. Woodford (1999), for example, proposes a rule in which the interest rate depends upon current and lagged values of the inflation rate as well as lagged interest rates in a similar model in which there is no inefficient supply shock. While his rule makes no reference to the output gap, it is dependent upon the driving process of the efficient rate of interest.

Equation (30) indicates that to implement the optimal plan, the central bank should relate the interest rate positively to fluctuations in current inflation, in *changes* of the output gap, and in lagged interest rates. While it is doubtful that the policymaker knows the current level of the output gap with great accuracy, the *change* in the output gap may be known with greater precision. For example, Orphanides (2003) shows that subsequent revisions of U.S. output gap estimates have been quite large (sometimes as large as 5.6 percentage points), while revisions of estimates of the quarterly change in the output gap have been much smaller.

Note finally that the interest rate should not only be inertial in the sense of being positively related to past values of the interest rate, it should be *super-inertial*, as the lagged polynomial for the interest rate in (30)

$$1 - \left(1 + \frac{\kappa}{\beta\sigma} + \beta^{-1}\right) \mathrm{L} + \beta^{-1}\mathrm{L}^2 = (1 - z_1\mathrm{L})(1 - z_2\mathrm{L})$$

involves a root $z_1 > 1$ while the other root $z_2 \in (0, 1)$. A reaction greater than one of the interest rate to its lagged value has initially been found by Rotemberg and Woodford (1999) to be a desirable

[20] The eigenvalues of this system are the same as the eigenvalues of M in (27) plus one eigenvalue equal to zero. As there is one predetermined variable more than in (27), this system yields a unique bounded equilibrium, if it exists.

[21] See Giannoni and Woodford (2002, 2003) for the characterization of policy rules which are optimal regardless of the shock processes in the context of more general models.

feature of a good policy rule in their econometric model with optimizing agents. As explained further in Woodford (2003), it is precisely such a *super-inertial* policy rule that the policymaker should follow to bring about the optimal responses to shocks when economic agents are forward-looking. Because of a root larger than one, the optimal policy requires an explosively growing response of the interest rate to deviations of inflation and the output gap from target.[22]

This is illustrated in Figure 3 which displays the response of the interest rate to a sustained 1 percent deviation in inflation (upper panel) or the output gap (lower panel) from target. In each panel, the solid line represents the optimal response in the baseline case. The corresponding coefficients of the optimal policy rule are reported in the upper panel of Table 2 (lines indicated by ψ^0). For comparison, the lower panel of Table 2 reports the coefficients derived from Judd and Rudebusch's (1998) estimation of actual Fed reaction functions between 1987:3 and 1997:4.[23] Table 2 reveals that the estimated historical rule in the baseline case involves only slightly smaller responses to fluctuations in inflation and the output gap than the optimal rule. However the estimated response to lagged values of the interest rate is sensibly smaller that the optimal one. As a result, the estimated historical rule involves a non-explosive response of the interest rate to a sustained deviation in inflation or the output gap, represented by the dashed-dotted lines in Figure 3.

Table 2 also reports the loss $\mathrm{E}[L_0]$ along with the following measure of variability

$$\mathrm{V}[z] \equiv \mathrm{E}\left\{\mathrm{E}_0\left[(1-\beta)\sum_{t=0}^{\infty}\beta^t z_t^2\right]\right\}$$

of the endogenous variables π, x and $\hat{\imath}$, for the various policy rules and parameter configurations. The statistic $\mathrm{V}[z]$ determines the contribution of each endogenous variable to the loss $\mathrm{E}[L_0]$, as the latter measure is a weighted sum of $\mathrm{V}[\pi]$, $\mathrm{V}[x]$, and $\mathrm{V}[\hat{\imath}]$ with weights corresponding to those of the loss function (18).[24] The lines of Table 2 indicated by θ^0 report statistics evaluated using the

[22] Interestingly, the rules preferred by Levin *et al.* (1999), while not super-inertial, involve a coefficient near unity on the lagged interest rate.

[23] The estimated historical policy rule refers to regression A for the Greenspan period in Judd and Rudebusch (1998).

[24] All statistics in Table 2 are expressed in annual terms. The statistics $\mathrm{V}[\pi]$, $\mathrm{V}[\hat{\imath}]$, and $\mathrm{E}[L_0]$ are therefore multiplied by 16. Furthermore, the weight λ_x reported in Table 1 is also multiplied by 16 in order to represent the weight attributed to output gap variability (in annual terms) relative to the variability of annualized inflation and of the annualized interest rate.

baseline parametrization. This table indicates to what extent the optimal rule results in a lower loss than the estimated historical rule.

While optimal policy would involve an explosive behavior of the interest rate in the face of a *sustained* deviation of inflation or the output gap, such a policy is perfectly consistent with a *stationary* rational expectations equilibrium, and a low variability of the interest rate *in equilibrium*. (In Table 2, $\mathrm{V}[\hat{\imath}]$ is always smaller when the interest rate is set according to the optimal flexible rule, than when it is set according to the estimated historical rule.) In fact, the interest rate does not explode in equilibrium because (as appears clearly in Figures 1 and 2) the current and expected future optimal levels of the interest rate are sufficient to counteract the effects of an initial deviation in inflation and the output gap by generating subsequent deviations with the opposite sign of these variables.

4.3 Robust Optimal Policy in the Presence of Parameter Uncertainty

So far we have assumed that all parameters are known with certainty. We now determine the robust optimal policy rule that obtains when the policymaker faces uncertainty about the three structural parameters σ, κ, and ω, the degrees of serial correlation of the exogenous perturbations ρ_δ, ρ_ε, ρ_μ, and the parameter ν describing the importance of the inefficient supply shocks. We consider uncertainty about the parameter vector $\theta = \left[\sigma, \kappa, \omega, \rho_\delta, \rho_\varepsilon, \rho_\mu, \nu\right]'$ within a specified set $\Theta = \left[\underline{\theta}, \bar{\theta}\right]$, where the extent of uncertainty for σ, κ, and ω is given by the approximate 95% confidence intervals mentioned in section 3.3, and reported in Table 1.[25] We assume that the coefficients of autocorrelation lie in the $[0, 0.8]$ interval,[26] and we allow ν to lie anywhere in the $[0, 1]$ interval. We assume that the preference parameters of the policymaker, λ_x and λ_i, are known to the policymakers, and are kept fixed regardless of the values of the structural parameters, for simplicity and because we view it as plausible that policymakers know their preferences despite the

[25] While the elasticities σ and ω are directly related to the underlying microeconomic foundations of the model, the parameter κ is a composite of structural parameters such as the degree of price rigidity α, the elasticity of substitution across goods. Since these parameters matter for welfare only through their effect on κ, and since the same aggregate supply equation can be derived with alternative microeconomic foundations, we find it more general to consider uncertainty about the parameter κ rather than some of its components.

[26] Alternative intervals yield similar results. It is important for our methodology, however, that the interval of uncertainty be closed.

fact that they don't know the true model.[27]

In the previous section, we have characterized the optimal equilibrium for any given parameter vector θ. Following the solution method summarized in section 2, we need to determine the candidate worst-case parameter vector, i.e., the parameter vector θ^* that obtains in the candidate minmax equilibrium. Once the worst-case parameter vector is identified, it will be straightforward to determine the robust optimal rule.

4.3.1 Worst-Case Parameter Vector and Minmax Equilibrium

We now determine the worst-case parameter vector θ^*, i.e., the vector θ that maximizes the loss function $\hat{L}\left(q^*\left(\theta\right)\right)$ on the constraint set Θ, or in other words the vector determining the minmax equilibrium $q^*\left(\theta^*\right)$. For any parameter vector $\theta \in \Theta$, the structural equations (14), (15) and the first-order conditions (23) – (25) can be written in matrix form, as in (27), and standard methods can be applied to get the solution of the form (28) – (29). As in the certainty case, the bounded solution is unique if one exists. Equations (28) – (29) can then be used to compute the loss $\hat{L}\left(q^*\left(\theta\right)\right)$. Maximizing this loss function with respect to $\theta \in \Theta$, we obtain:

$$\begin{aligned} \theta^* &= \left[\underline{\sigma}, \bar{\kappa}, \underline{\omega}, \bar{\rho}_\delta, \rho_\varepsilon, \bar{\rho}_\mu, \bar{\nu}\right]' \\ &= \left[0.0915, 0.0308, 0.2837, 0.8, -, 0.8, 1\right]. \end{aligned} \tag{32}$$

Note that ρ_ε may take any value in the allowed interval $[0, 1]$, since the loss is maximized when $\nu^* = 1$, i.e., when there are no efficient supply shocks. We performed the maximization of $\hat{L}\left(q^*\left(\theta\right)\right)$ numerically starting from a large number of different initial values for θ, including values close to the boundaries of the set Θ. There are a few local maxima, but none of them yields a loss higher than the one implied by the parameter vector reported above.

The minmax equilibrium involves very persistent shocks in the more general model of the previous sections. Indeed, the worst-case values for the coefficients of serial correlations are equal

[27] Alternatively, if the policymaker's loss function is interpreted strictly as a second-order approximation to the expected utility of the representative household in the underlying microeconomic model, and there is no uncertainty about how the model relates to the policymaker's preferences, the parameters λ_x and λ_i are related to the model's structural parameters. As a result uncertainty about the underlying model parameters should translate into uncertainty about λ_x and λ_i. Kurozumi (2003) and Walsh (2005) consider such a case.

to the upper bound 0.8. Note that this is true even though the variance-covariance matrix of the shocks (20) is given, and independent of ρ_δ, ρ_ε, and ρ_μ.

To get some intuition about the worst-case parameter values for σ, κ, and ω, we consider temporarily a simple case in which monetary policy is assumed to be non-inertial — i.e., it does not depend on lagged variables — and all shocks are i.i.d. In this case, all future variables are expected to remain at steady state in equilibrium ($\mathrm{E}_t x_{t+1} = \mathrm{E}_t \pi_{t+1} = 0$), so that the two structural equations (14) and (15) reduce to

$$\hat{\imath}_t = -\sigma x_t + \frac{\omega}{\omega+\sigma}\delta_t + \frac{\sigma}{\omega+\sigma}\varepsilon_t \tag{33}$$

$$\pi_t = \kappa\left(x_t + \frac{1}{\omega+\sigma}\mu_t\right). \tag{34}$$

These two equations are represented by respectively the lines IS and AS in Figures 4a to 4c.

Figure 4a represents the effects of a unit exogenous increase in δ_t. In the case in which the parameters are known with certainty, the IS curve shifts vertically from $IS(\sigma_0, \omega_0)$ to $IS'(\sigma_0, \omega_0)$, by an amount $\frac{\omega_0}{\omega_0+\sigma_0}$. The policymaker faces a trade-off between the stabilization of inflation and the output gap on one hand, and the interest rate on the other hand. He could completely stabilize inflation and the output gap by raising the interest rate by $\frac{\omega_0}{\omega_0+\sigma_0}$. Such a policy is however not optimal as the policymaker also cares about fluctuations in the interest rate (see loss function (18)). He acts optimally by increasing the interest rate to some level i_t^0, and letting the output gap increase to x_t^0. In the lower panel of the figure, inflation rises to π_t^0. In the presence of parameter uncertainty and a demand shock, the worst case is obtained when σ is as low as possible, and ω is as high as possible, so that $\sigma^* = \underline{\sigma}$ and $\omega^* = \bar{\omega}$. As (33) reveals, it is in this case that a given increase increase in δ_t results in the highest possible upward shift of the IS curve. Furthermore, as this implies a flatter IS curve, the output gap increases by more, for given nominal interest rate. On the supply side, the worst value for κ is obtained when $\kappa^* = \bar{\kappa}$, so that any given change in the output gap is associated with a large change in inflation. It follows that the policymaker who seeks to minimize the loss in the worst-case parameter configuration optimally sets the interest rate above i_t^0 in the presence of uncertainty, in order to contain the increase both in the output gap and in inflation.[28]

[28] The fact that the nominal interest rate rises above i_t^0 in the presence of uncertainty does not necessarily mean that

Figure 4b shows that the effects of a unit exogenous increase in ε_t are qualitatively similar to those of a shock to δ_t when parameters are known with certainty. Note that it is the IS and not the AS schedule that shifts following an efficient supply shock. The IS schedule shifts upwards to IS' by an amount $\frac{\sigma_0}{\omega_0+\sigma_0}$. However, in the presence of parameter uncertainty, the worst-case value for ω is $\underline{\omega}$, because it generates the largest shift of the IS schedule. As before, the worst-case value for κ is $\bar{\kappa}$, so that a given non-zero output gap results in the largest change in inflation. It is however not trivial to determine a priori the worst-case value for σ. Whether σ^* is $\underline{\sigma}$, $\bar{\sigma}$, or any value in between depends on the parameter values. For the calibration of Table 1, it is $\bar{\sigma}$, as this is the value which is responsible for the largest upward shift of IS, even though it implies a steeper IS schedule.

Figure 4c illustrates the effects of an inefficient supply shock, which shifts the AS curve from $AS(\kappa_0)$ up to $AS'(\kappa_0)$, in the absence of parameter uncertainty. The policymaker faces a trade-off between the stabilization of inflation on one hand, and the output gap on the other hand. He acts optimally by raising the interest rate to some level i_t^0, so that the output gap decreases to x_t^0, and inflation increases to π_t^0. In the presence of parameter uncertainty, the worst-case slope of the AS curve is again obtained when $\kappa^* = \bar{\kappa}$, and the horizontal shift of the AS curve is largest when $\sigma^* = \underline{\sigma}$, and $\omega^* = \underline{\omega}$. In the upper panel, however, we notice that the worst slope of the IS curve obtains when $\sigma^* = \bar{\sigma}$, so that the policymaker needs to increase the interest rate by more, to obtain a given change in the output gap.

To summarize, while the worst-case value for κ is $\bar{\kappa}$ regardless of the shock considered, the worst-case parameter values for σ and ω depend in general on the parametrization of the model, and the relative importance of the disturbances. For the parametrization summarized in Table 1, the worst-case parameter configuration involves $\sigma^* = \underline{\sigma}$ and $\omega^* = \underline{\omega}$ in the model considered.

4.3.2 Robust Optimal Policy Rule

Following the solution procedure reviewed in section 2, once the worst-case parameter vector θ^* has been determined, we characterize the robust optimal policy rule simply by looking at the best response $\psi^*(\theta^*)$ to this vector, assuming that a global NE does exist. As (31) is the best response

the policy rule involves larger responses to variables such as inflation and the output gap, since the latter variables may increase too. However, we know from (30) that the robust optimal rule requires larger responses to fluctuations in inflation and the output gap when the worst-case values for σ and κ are respectively $\underline{\sigma}$ and $\bar{\kappa}$.

to any given θ, the candidate *robust optimal policy rule*, $\psi^* = \psi^*(\theta^*)$ satisfies

$$\hat{\imath}_t = \frac{\bar{\kappa}}{\lambda_i \underline{\sigma}} \pi_t + \frac{\lambda_x}{\lambda_i \underline{\sigma}} (x_t - x_{t-1}) + \left(1 + \frac{\bar{\kappa}}{\beta \underline{\sigma}} + \beta^{-1}\right) \hat{\imath}_{t-1} - \beta^{-1} \hat{\imath}_{t-2}. \tag{35}$$

The couple (ψ^*, θ^*) constitutes at least a local NE, and the policy rule satisfying (35) is a *robust optimal policy rule*, provided that a global NE exists.

In accordance with step 4 of our solution procedure, we verify that the couple (ψ^*, θ^*) is also a global NE. We do so by maximizing the loss function $L(\psi^*, \theta) \equiv \mathrm{E}[L_0(q(\psi^*, \theta))]$ numerically with respect to $\theta \in \Theta$. This loss function results from the equilibrium obtained by combining the structural equations (14), (15) for any given θ, and the candidate robust optimal policy rule ψ^* satisfying (35). We repeated the maximization many times, starting each time from a different initial values for θ. We obtained again that the vector θ^* defined in (32) maximizes this loss function. It follows that by choosing θ^*, malevolent Nature best-responds to the policy rule ψ^*, and that (ψ^*, θ^*) is indeed a global NE. Because θ^* is the only equilibrium parameter vector, and $\psi^* \equiv \psi^*(\theta^*)$ is the unique best response to θ^* in the set $\tilde{\Psi}$, the profile (ψ^*, θ^*) constitutes the unique global NE, and ψ^* is the *unique* robust optimal policy rule in the class $\tilde{\Psi}$.

Denoting by $\sigma_0 \in (\underline{\sigma}, \bar{\sigma})$ and $\kappa_0 \in (\underline{\kappa}, \bar{\kappa})$ the parameter values in the absence of model uncertainty, and noting that $\underline{\sigma} < \sigma_0$ and $\kappa_0 < \bar{\kappa}$, one observes, by comparing (30) and (35) that the policymaker reacts *more* strongly to perturbations to inflation, changes in the output gap, and the lagged interest rate than is the case in the absence of uncertainty. To give a sense of the magnitude of optimal policy coefficients, we report in the middle panel of Table 2 the robust optimal rule (35) (lines indicated by ψ^*), in addition to the optimal policy rule (30) in the certainty case (lines ψ^0). The lagged polynomial for the interest rate can be written as

$$\begin{aligned} \text{Baseline} \quad &: \quad 1 - 2.163\mathrm{L} + 1.010\mathrm{L}^2 = (1 - 1.481\mathrm{L})(1 - .682\mathrm{L}) \\ \text{Robust} \quad &: \quad 1 - 2.350\mathrm{L} + 1.010\mathrm{L}^2 = (1 - 1.784\mathrm{L})(1 - .566\mathrm{L}). \end{aligned}$$

As the larger root is even greater in the presence of uncertainty, the interest rate is *super-inertial* to an even greater extent when the central bank follows the robust optimal rule. As illustrated

in Figure 3, the robust optimal rule (dashed line) involves (i) a *larger* response to a sustained increase in inflation, at every time, (ii) a *larger* response to a sustained increase in the output gap, at every time, and (iii) a *faster* asymptotic rate of explosion of both of these responses (thus a greater degree to which the policy rule is *super-inertial*). But again, the presence of a root larger than one is consistent with a stationary rational expectations equilibrium. This is illustrated for instance by the dashed lines in Figures 1 and 2, which represent the impulse responses of the endogenous variables when the policymaker follows the robust optimal policy rule (but when the true parameters are those of the baseline calibration).

4.4 Sensitivity of Robust Policy to Alternative Assumptions

We just argued that the robust optimal policy rule involves larger responses of the interest rate to endogenous variables. Equation (35) indicates however that the robust optimal rule depends critically on the worst-case values for σ and κ. While the parameters $\omega, \rho_\delta, \rho_\varepsilon, \rho_\mu, \nu$ don't enter directly the policy rule (35), they may still have an indirect effect by affecting the worst-case values for σ and κ. To get a sense of how robust our conclusions are, we consider alternative assumptions about the uncertainty surrounding the structural parameters.

4.4.1 Robust Policy with Little Uncertainty

As discussed in section 4.3.1, the worst-case value for κ remains the upper bound $\bar{\kappa}$ regardless of the importance of the shocks considered, and the values of the other structural parameters — at least when monetary policy is non-inertial. This induces the policymaker to let the interest rate react more strongly to fluctuations in inflation and lagged interest rates, in the presence of parameter uncertainty. The story is potentially more complicated for the parameter σ. For the calibration summarized in Table 1, though, the worst-case value is $\underline{\sigma}$. Figure 5 indicates that $\bar{\kappa}$ and $\underline{\sigma}$ remain the worst-case values for κ and σ, under alternative assumptions about ω and the degree of serial correlation of the perturbations, ρ, when ν is maintained at its worst-case value $\nu^* = 1$. This figure represents contour plots of the loss criterion $\mathrm{E}[L_0]$ as a function of the parameters κ and σ, when monetary policy is conducted according to the robust optimal rule (35). These contour plots are produced for various values of ω and ρ, and we report the cases in which ω is respectively $\underline{\omega}$, the

baseline value ω_0, and $\bar{\omega}$, and ρ is respectively 0, 0.35, and 0.8. The star in each plot indicates the baseline values for σ and κ. The figure reveals that the worst-case couple (κ^*, σ^*) — indicated by a circled star — is in each case in the lower right corner, i.e., $(\bar{\kappa}, \underline{\sigma})$. Similar results are obtained for alternative values for ω and ρ, and for alternative values of ν, provided that ν is larger than a critical value around 0.5. This suggests that the robust optimal rule is not affected by alternative assumptions about ω and ρ, as long as the maximum value allowed for ν, i.e., $\bar{\nu}$ is large enough.

However, if $\bar{\nu}$ is sufficiently small — so that most supply shocks are *known* to be mostly efficient supply shocks ε_t — then the worst-case value for σ may be $\bar{\sigma}$, as argued in subsection 4.3.1. For instance, if ν is constrained to lie in the interval $[0, 0.3]$, then the worst-case value for ν turns out to be its lower bound 0. It follows that the worst case value for σ is $\bar{\sigma}$. In this case, the response to fluctuations in the output gap would be smaller in the presence of uncertainty than in the certainty case, as $\lambda_x \left(\lambda_i \bar{\sigma}\right)^{-1} < \lambda_x \left(\lambda_i \sigma_0\right)^{-1}$. Whether the robust rule responds more to fluctuations in inflation and in the lagged interest rate depends in the end on the amount of uncertainty about κ relative to the one about σ. In our numerical example, if the worst case value for σ is $\bar{\sigma}$, it appears from (35) that the response coefficients to fluctuations in inflation and the lagged interest rate are also smaller in the presence of uncertainty, since $\bar{\kappa}/\bar{\sigma} = 0.0308/0.2227 < 0.0238/0.1571 = \kappa_0/\sigma_0$.

To summarize, the robust optimal policy rule (35) does not depend on ω, ν and the coefficients of serial correlation of the shocks, provided that the uncertainty about ν is such that $\bar{\nu}$ is large enough. However, if $\bar{\nu}$ is constrained to be sufficiently small, then the worst-case value for σ is $\bar{\sigma}$, and the robust optimal rule is less aggressive than the optimal rule absent parameter uncertainty.

4.4.2 Robust Policy with Large Uncertainty about σ, κ, ω

The analysis has so far assumed that the uncertainty about the critical parameters σ, κ, ω, is bounded by an estimate of their respective 95% confidence intervals. The degree of parameter uncertainty assumed might however be considered as relatively small. To check further the robustness of our results, we consider an alternative calibration of the amount of uncertainty by allowing the

intervals to be considerably larger. Specifically, we assume

$$[\underline{\sigma}, \bar{\sigma}] = [0.05, 1] \tag{36}$$

$$[\underline{\kappa}, \bar{\kappa}] = [0.01, 0.5] \tag{37}$$

$$[\underline{\omega}, \bar{\omega}] = [0.1, 1]\,. \tag{38}$$

We repeat the steps laid down in section 2. For this alternative calibration, we find again that the worst-case parameter configuration involves $\sigma^* = \underline{\sigma}$, $\kappa^* = \bar{\kappa}$ and $\omega^* = \underline{\omega}$. The robust optimal policy rule is thus again of the form (35).

By how much the policymaker should let the interest rate respond to fluctuations of inflation, the output gap and past values of the interest rate actually depends on the particular degree of uncertainty assumed, in particular on the values attributed to $\bar{\kappa}$ and $\underline{\sigma}$. In any case, the qualitative result remains the same: in the face of uncertainty about the structural parameters, the robust optimal rule involves larger responses of the interest rate to fluctuations in inflation, output gap and lagged interest rate, than is the case in the absence of uncertainty, and that response is larger the higher $\bar{\kappa}$ and the lower $\underline{\sigma}$.

4.5 Welfare Comparisons

In the model considered, the worst-case parameter vector lies at the boundary of the relevant set.[29] Such a parameter configuration may thus be viewed as very unlikely, compared to other parameter configurations. One may thus not find it attractive to adopt a policy rule that is optimal only in some unlikely situation. It is important to remember, however, that the robust policy rule is not optimal for any arbitrary unlikely parameter vector. By being optimal for the worst-case situation, it guarantees that the losses encountered will not be "too" large.

To get a sense of the benefits and costs of adopting a robust optimal rule, we compare in Figures 6-10 its welfare implications assuming different parameter situations. Figure 6 compares the performance of the robust optimal rule ψ^* and the rule that is optimal in the absence of parameter

[29] While this does not always need to be the case in the context of parameter uncertainty (see, e.g., Giannoni, 2002), analyses of unstructured uncertainty *a la* Hansen and Sargent (2005) do involve a worst-case scenario on the boundary of the set of relevant models.

uncertainty (ψ^0) both in the baseline case $\left(\theta^0\right)$ and in the worst-case parameter configuration (θ^*). Each panel plots the loss as a function of one parameter, keeping the remaining parameters either at their baseline or their worst-case value. The figure reveals that the welfare losses are much larger for the worst-case parametrization than in the benchmark case, regardless of the policy rule adopted. While both rules result in almost identical losses in the baseline calibration (θ^0), it appears that the robust rule performs marginally better than the baseline rule when the parameters are close to their worst-case value.[30] Note however from Table 2 that both ψ^0 and ψ^* perform significantly better than the estimated historical rule.

Even though the overall welfare losses are broadly similar with the robust optimal rule ψ^* and the optimal rule absent model uncertainty (ψ^0), Figures 7-9 show that the policy rules have very different implications for the volatility of inflation, the output gap and the nominal interest rate. The robust optimal policy brings about a lower variability of inflation than the rule ψ^0, but it is responsible for a larger volatility of the output gap and the nominal interest rate than the rule ψ^0. Thus, in the model considered, the robust rule is one that guards mostly against inflation volatility.

Figure 6 displayed modest welfare gain from the robust optimal rule. This is however due to the relatively small degree of uncertainty is assumed about the parameters κ and σ. If instead, we assume that σ, κ, ω lie in the larger intervals given by (36) – (38), we observe from Figure 10 that the robust optimal rule described in the subsection 4.4.2 achieves much lower welfare losses than the rule ψ^0 for parameter configurations not even close from the worst-case situation. In contrast, around the baseline calibration, the robust optimal rule's welfare performance is only marginally worse than the one of ψ^0. It should be clear from this example, that when the degree of parameter uncertainty is sufficiently large, the benefits of following the robust rule dominate the costs.

5 Conclusion

In this paper, we have characterized a robust optimal policy rule in a simple forward-looking model, when the policymaker faces uncertainty about the parameters of the structural model and

[30] The figure also reveals that if all elements of θ with the exception of one, reach their worst-case values at the respective values in θ^*, then the remaining element of θ also maximizes the loss at the corresponding value in θ^*. This illustrates the fact that θ^* is a best response on the part of malevolent Nature to the robust rule ψ^*.

the nature of the shock processes. We have derived the structural model from first principles to determine precisely how the exogenous perturbations are transmitted to the endogenous variables.

The optimal policy rule considered here has a number of advantages with respect to simpler policy rules such as the Taylor rule. First, as it implements the optimal plan in the absence of parameter uncertainty, it achieves the lowest possible loss, and hence performs better than restricted policy rules.[31] Second, the analytical characterization of the optimal rule allows us to identify to what extent the policy rule is sensitive to particular parameters. While the optimal Taylor rule derived in Giannoni (2002) depends critically on the characteristics of the exogenous shock processes, the optimal rule proposed here does not depend them, in the absence of parameter uncertainty. The invariance to various specifications of the shock processes is an attractive feature of the optimal rule, especially when exogenous disturbances cannot be observed directly. The *robust* optimal rule depends however indirectly on the assumptions about the shock processes to the extent that they affect the worst-case parameter configuration, but this turns out having little impact. Another interesting feature of the optimal rule is that it is *super-inertial*, i.e., it involves response coefficients to lagged interest rates that are larger than one. As first shown in Rotemberg and Woodford (1999), and Woodford (1999), this feature of monetary policy allows the central bank to affect the private sector's expectations appropriately.

We have shown that in the presence of parameter uncertainty, the robust policy rule implies that the interest rate responds in general more strongly to fluctuations in inflation, to changes in the output gap, and is super-inertial to an even greater extent, than is the case in the absence of uncertainty. This result holds quite generally in the model considered, provided that there is enough uncertainty in particular about the type of supply shocks. Even if he responds more strongly, in the presence of parameter uncertainty, the policymaker is *cautious* in our framework. In fact he is even more cautious than in Brainard's model, as he cares very much about worst-case situations. We presented an example in which "being cautious" does not necessarily mean "to do less".

The robust optimal rule derived here is designed to be robust against a specific type of uncertainty, namely uncertainty about some key parameters of the model, and the results obtained here are specific to the environment considered. One may however be uncertain about many other

[31] For a comparison of the performance of both rules in the absence of parameter uncertainty, see Giannoni (2000).

aspects of the model such as the formation of expectations, the variables entering each equation, the number of lags of each variable, and so on. Some of these types of uncertainty can be studied in the context of parameter uncertainty, and the approach described in section 2 can still be applied in such cases. This approach has the advantage of being transparent, allowing a simple characterization of the robust optimal rule, and providing an intuitive understanding of the results obtained. However if the uncertainty considered involves disjoint models or, in significantly larger models, uncertainty about a very large number of parameters, the proposed approach may not be feasible. In this case, alternative methods imposing less structure on the uncertainty, or methods based on model averaging may be better suited.

References

[1] Amato, Jeffery D., and Thomas Laubach (2003), "Estimation and control of an optimization-based model with sticky prices and wages," *Journal of Economic Dynamics and Control* 27(7), 1181-1215.

[2] Aoki, Kosuke (2003) "On the optimal monetary policy response to noisy indicators," *Journal of Monetary Economics* 50(3): 501-523.

[3] Ball, Laurence (1999) "Policy rules for open economies," in John B. Taylor (ed.), *Monetary Policy Rules,* pp. 127-144, Chicago: University of Chicago Press.

[4] Blanchard, Olivier J., and Charles Kahn (1980) "The solution of linear difference equations under rational expectations," *Econometrica* 48: 1305-1311.

[5] Blinder, Alan. S. (1998) *Central Banking in Theory and Practice,* Cambridge, MA: MIT Press.

[6] Brainard, William (1967) "Uncertainty and the effectiveness of policy," *American Economic Review* 57(2): 411-425.

[7] Brock, William A., Steven N. Durlauf, and Kenneth D. West (2003), "Policy evaluation in uncertain economic environments," NBER working paper 10025.

[8] —— ——, —— ——, and —— —— (2004) "Model uncertainty and policy evaluation: Some theory and empirics," NBER working paper 10916.

[9] Calvo, Guillermo A. (1983), "Staggered prices in a utility-maximizing framework," *Journal of Monetary Economics* 12(3): 383-398.

[10] Chow, Gregory C. (1975) *Analysis and Control of Dynamic Economic Systems,* New York, NY: John Wiley & Sons.

[11] Clarida, Richard, Jordi Galí, and Mark Gertler (1999) "The science of monetary policy: a New Keynesian perspective," *Journal of Economic Literature* 37(4): 1661-1707.

[12] —— ——, —— ——, and —— —— (2000) "Monetary policy rules and macroeconomic stability: evidence and some theory," *Quarterly Journal of Economics* 115, 147-180.

[13] Christiano, Lawrence J., and Christopher Gust (1999) "Comment," in John B. Taylor (ed.), *Monetary Policy Rules,* pp. 299-316, Chicago: University of Chicago Press.

[14] Dixit, Avinash K., and Joseph E. Stiglitz (1977) "Monopolistic competition and optimum product diversity," *American Economic Review* 67: 297-308.

[15] Estrella, Arturo, and Frederic S. Mishkin (1999) "Rethinking the role of NAIRU in monetary policy: implications of model formulation and uncertainty," in John B. Taylor (ed.), *Monetary Policy Rules,* pp. 405-430, Chicago: University of Chicago Press.

[16] Friedman, Milton (1969) "The optimum quantity of money," in Milton Friedman, *The Optimum Quantity of Money and Other Essays.* Chicago: Aldine.

[17] Giannoni, Marc P. (2000) "Optimal interest-rate rules in a forward-looking model and inflation stabilization versus price-level stabilization," manuscript, Columbia University.

[18] —— —— (2002) "Does model uncertainty justify caution? Robust optimal monetary policy in a forward-looking model," *Macroeconomic Dynamics* 6(1), 111-144.

[19] —— ——, and Michael Woodford (2002), "Optimal interest-rate rules: I. General theory," NBER working paper no. 9419.

[20] —— ——, and Michael Woodford (2003), "How forward-looking is optimal monetary policy?" *Journal of Money, Credit, and Banking* 35(6), 1425-1469.

[21] Gilboa, Itzhak and David Schmeidler (1989) "Maxmin expected utility with non-unique prior," *Journal of Mathematical Economics* 18: 141-153.

[22] Hall, Simon, Chris Salmon, Tony Yates, and Nicoletta Batini (1999) "Uncertainty and simple monetary policy rules: an illustration for the United Kingdom," Bank of England working paper no. 96.

[23] Hansen, Lars P., and Thomas J. Sargent (2000) "Wanting robustness in macroeconomics," manuscript, University of Chicago, New York University and Hoover Institution.

[24] —— ——, and —— —— (2003) "Robust control of forward-looking models," *Journal of Monetary Economics* 50(3): 581-604.

[25] —— ——, and —— —— (2005) *Robustness*, Princeton, Princeton University Press, forthcoming.

[26] Judd, John P., and Glenn D. Rudebusch (1998) "Taylor's rule and the Fed: 1970-1997," Federal Reserve Bank of San Francisco, *Economic Review* 3: 3-16.

[27] Kasa, Kenneth (2002) "Model uncertainty, robust policies, and the value of commitment," *Macroeconomic Dynamics* 6(1), 145-166.

[28] Kurozumi, Takushi (2003) "Optimal monetary policy under parameter uncertainty in a simple micro-founded model," manuscript, Carnegie Mellon University.

[29] Knight, Frank H. (1921) *Risk, Uncertainty and Profit*, Boston, New York, Houghton Mifflin Company.

[30] Levin, Andrew, Volker Wieland, and John C. Williams (1999) "Robustness of simple monetary policy rules under model uncertainty," in John B. Taylor (ed.), *Monetary Policy Rules,* pp. 263-299. Chicago: University of Chicago Press.

[31] —— ——, —— ——, and —— —— (2003) "The performance of forecast-based monetary policy rules under model uncertainty," *American Economic Review* 93(3), 622-645.

[32] Levin, Andrew, and John C. Williams (2003) "Robust monetary policy with competing reference models," *Journal of Monetary Economics* 50(5): 945-975.

[33] Martin, Ben, and Chris Salmon (1999) "Should uncertain monetary policy-makers do less?" Bank of England working paper no. 99.

[34] McCallum, Bennett T. (1988) "Robustness properties of a rule for monetary policy," *Carnegie-Rochester Conference Series on Public Policy* 29: 173-204.

[35] —— —— (1999) "Issues in the design of monetary policy rules," in John B. Taylor and Michael Woodford (eds.), *Handbook of Macroeconomics*, Vol. 1C. Amsterdam: North-Holland.

[36] Onatski, Alexei (2000a) "Minimax monetary policy: Comparison to Bayesian approach, worst cases, and exact minimax rules," manuscript, Harvard University.

[37] —— —— (2000b) "Robust monetary policy under model uncertainty: Incorporating rational expectations," manuscript, Harvard University.

[38] —— ——, and James H. Stock (2002) "Robust monetary policy under model uncertainty in a small model of the U.S. economy," *Macroeconomic Dynamics* 6(1), 85-110.

[39] —— ——, and Noah Williams (2003) "Modeling model uncertainty," *Journal of the European Economic Association* 1(5): 1087-1122.

[40] Orphanides, Athanasios (2003) "Monetary policy evaluation with noisy information," *Journal of Monetary Economics* 50(3): 605-631.

[41] Rotemberg, Julio J., and Michael Woodford (1997) "An optimization-based econometric framework for the evaluation of monetary policy," *NBER Macroeconomics Annual 1997,* pp. 297-346.

[42] —— ——, and —— —— (1999) "Interest rate rules in an estimated sticky price model," in John B. Taylor (ed.), *Monetary Policy Rules,* pp. 57-119. Chicago: University of Chicago Press.

[43] Rudebusch, Glenn D. (2001) "Is the Fed too timid? Monetary policy in an uncertain world," *Review of Economics and Statistics* 83(2), 203-217.

[44] Sack, Brian (2000) "Does the Fed act gradually? A VAR analysis," *Journal of Monetary Economics* (46)1: 229-256.

[45] Sargent, Thomas J. (1999) "Comment," in John B. Taylor (ed.), *Monetary Policy Rules,* pp. 144-154. Chicago: University of Chicago Press.

[46] Stock, James H. (1999) "Comment," in John B. Taylor (ed.), *Monetary Policy Rules,* pp. 253-259. Chicago: University of Chicago Press.

[47] Söderström, Ulf (2000) "Should central banks be more aggressive?" manuscript, Sveriges Riksbank and Stockholm School of Economics.

[48] —— —— (2002) "Monetary policy with uncertain parameters," *Scandinavian Journal of Economics* 104 (1), 125-145.

[49] Svensson, Lars E. O. (1999) "Inflation targeting: some extensions," *Scandinavian Journal of Economics* 101 (3): 337-361.

[50] Svensson, Lars E. O., and Michael Woodford (2003), "Indicator Variables for Optimal Policy," *Journal of Monetary Economics* 50, 691–720.

[51] —— ——, and —— —— (2004), "Indicator Variables for Optimal Policy under Asymmetric Information," *Journal of Economic Dynamics and Control* 28, 661–690.

[52] Taylor, John B. (1993) "Discretion versus policy rules in practice," *Carnegie-Rochester Conference Series on Public Policy* 39, 195-214.

[53] —— —— (ed.) (1999a) *Monetary Policy Rules.* Chicago: University of Chicago Press.

[54] —— —— (1999b) "Introduction," in John B. Taylor (ed.), *Monetary Policy Rules,* pp. 1-14. Chicago: University of Chicago Press.

[55] Tetlow, Robert J., and Peter 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, 911-949.

[56] von zur Muehlen, Peter (1982) "Activist vs. non-activist monetary policy: Optimal rules under extreme uncertainty," manuscript. Federal Reserve Board.

[57] Walsh, Carl E. (2003), "Implications of a changing economic structure for the strategy of monetary policy," *Monetary Policy Uncertainty: Adapting to a Changing Economy*, Federal Reserve Bank of Kansas City, Jackson Hole Symposium, pp. 297-348.

[58] —— —— (2005), "Endogenous objectives and the evaluation of targeting rules for monetary policy," forthcoming in *Journal of Monetary Economics*.

[59] Wieland, Volker (1998) "Monetary policy and uncertainty about the natural unemployment rate," FEDS discussion paper no. 1988-22, Federal Reserve Board.

[60] Woodford, Michael (1990) "The optimum quantity of money," in B.M. Friedman and F.H. Hahn (eds.), *Handbook of Monetary Economics*, Vol. II. Amsterdam: North-Holland.

[61] —— —— (1999) "Optimal monetary policy inertia" NBER working paper no. 7261.

[62] —— —— (2003), *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton, Princeton University Press.

Table 1: Calibrated Parameter Values

Structural parameters (benchmark case)

	β	σ	κ	ω
estimate:	0.99	0.1571	0.0238	0.4729
std. error:	–	0.0328	0.0035	0.0946
lower bound:	–	0.0915	0.0168	0.2837
upper bound:	–	0.2227	0.0308	0.6621

Shock processes

Variance-covariance matrix

of $u_t \equiv [\delta_t, \varepsilon_t, \mu_t]'$ when $\nu = 0.5$

3.0150	1.6058	14.1131
	43.9248	39.1573
		122.9095

serial correlation: $\rho_\delta = \rho_e = \rho_\mu = 0.35$

Loss function

λ_x	λ_i
0.048	0.236

Table 2: Policy Rules and Statistics

		Coefficients of policy rule					Statistics			
		π_t	x_t	x_{t-1}	i_{t-1}	i_{t-2}	V[π]	V[x]	V[i]	E[L_0]
Optimal Rules										
ν = 0	(ψ^0,θ^0)	0.641	0.325	-0.325	2.163	-1.010	0.130	10.599	1.921	1.097
	(ψ^0,θ^*)	0.641	0.325	-0.325	2.163	-1.010	0.408	9.488	5.838	2.247
	(ψ^*,θ^0)	1.424	0.559	-0.559	2.350	-1.010	0.126	7.334	2.806	1.144
	(ψ^*,θ^*)	1.424	0.559	-0.559	2.350	-1.010	0.366	5.325	6.635	2.192
ν = .5	(ψ^0,θ^0)	0.641	0.325	-0.325	2.163	-1.010	0.213	4.435	0.718	0.597
	(ψ^0,θ^*)	0.641	0.325	-0.325	2.163	-1.010	0.790	24.659	2.116	2.482
	(ψ^*,θ^0)	1.424	0.559	-0.559	2.350	-1.010	0.182	3.831	1.081	0.622
	(ψ^*,θ^*)	1.424	0.559	-0.559	2.350	-1.010	0.592	26.086	2.439	2.429
ν = 1	(ψ^0,θ^0)	0.641	0.325	-0.325	2.163	-1.010	0.569	5.759	0.257	0.908
	(ψ^0,θ^*)	0.641	0.325	-0.325	2.163	-1.010	2.431	88.093	0.724	6.859
	(ψ^*,θ^0)	1.424	0.559	-0.559	2.350	-1.010	0.490	7.057	0.415	0.929
	(ψ^*,θ^*)	1.424	0.559	-0.559	2.350	-1.010	1.833	97.981	0.848	6.769
Estimated Historical										
ν = 0	(ψ,θ^0)	0.424	0.297	-0.032	1.160	-0.430	0.079	11.852	2.952	1.349
	(ψ,θ^*)	0.424	0.297	-0.032	1.160	-0.430	1.284	14.891	14.184	5.357
ν = .5	(ψ,θ^0)	0.424	0.297	-0.032	1.160	-0.430	0.465	3.737	1.504	1.001
	(ψ,θ^*)	0.424	0.297	-0.032	1.160	-0.430	11.782	9.122	17.341	16.322
ν = 1	(ψ,θ^0)	0.424	0.297	-0.032	1.160	-0.430	1.363	1.469	0.959	1.661
	(ψ,θ^*)	0.424	0.297	-0.032	1.160	-0.430	36.041	22.245	27.291	43.568

Note: The estimated historical rule refers to Judd and Rudebusch (1998).

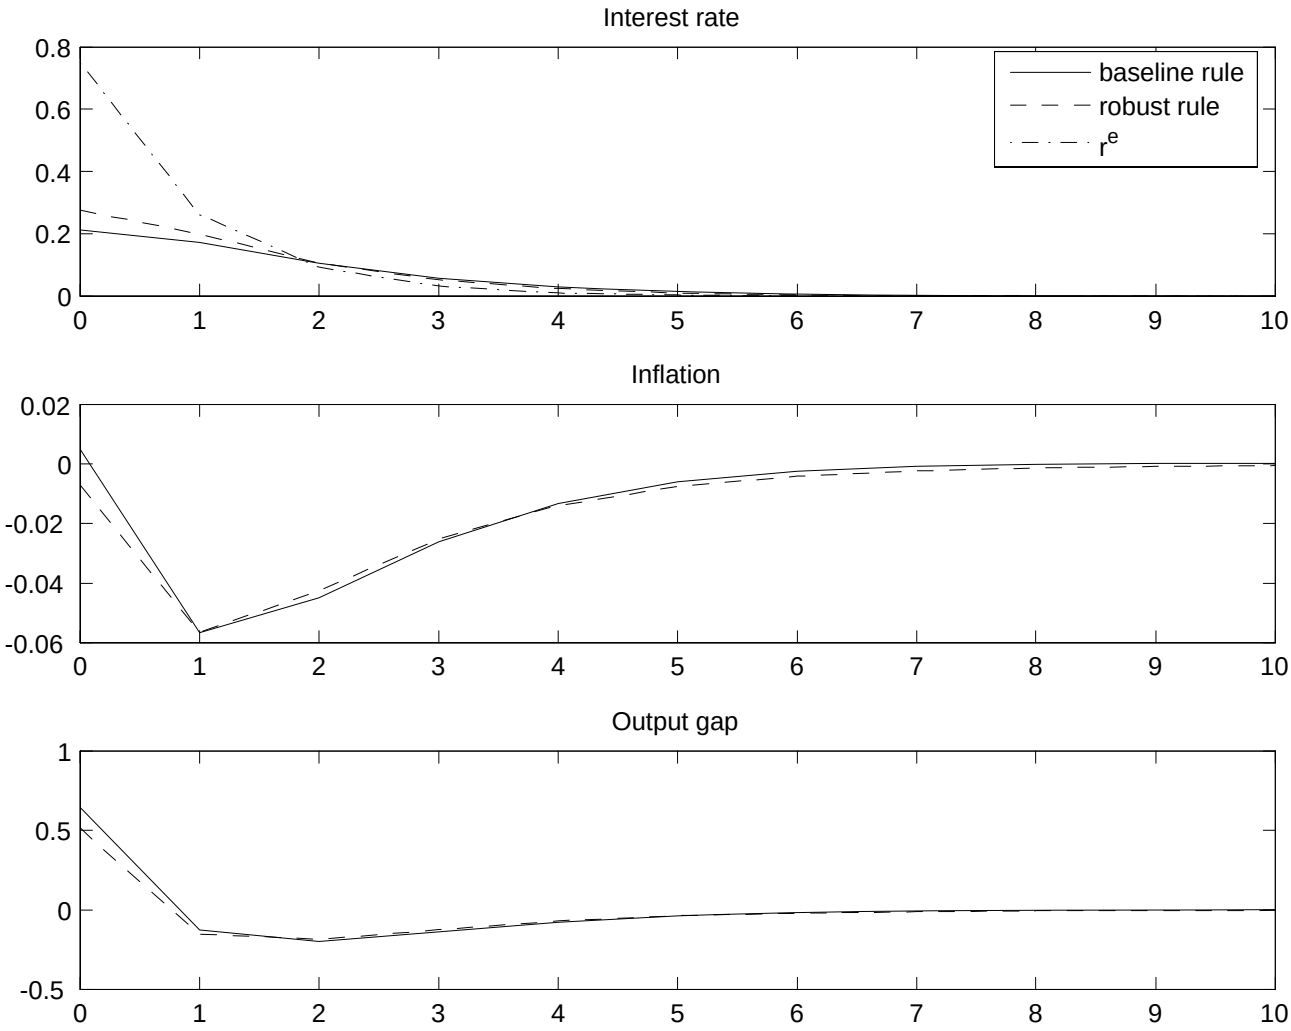

Figure 1: Impulse responses to a shock δ ($\rho_\delta = 0.35$).

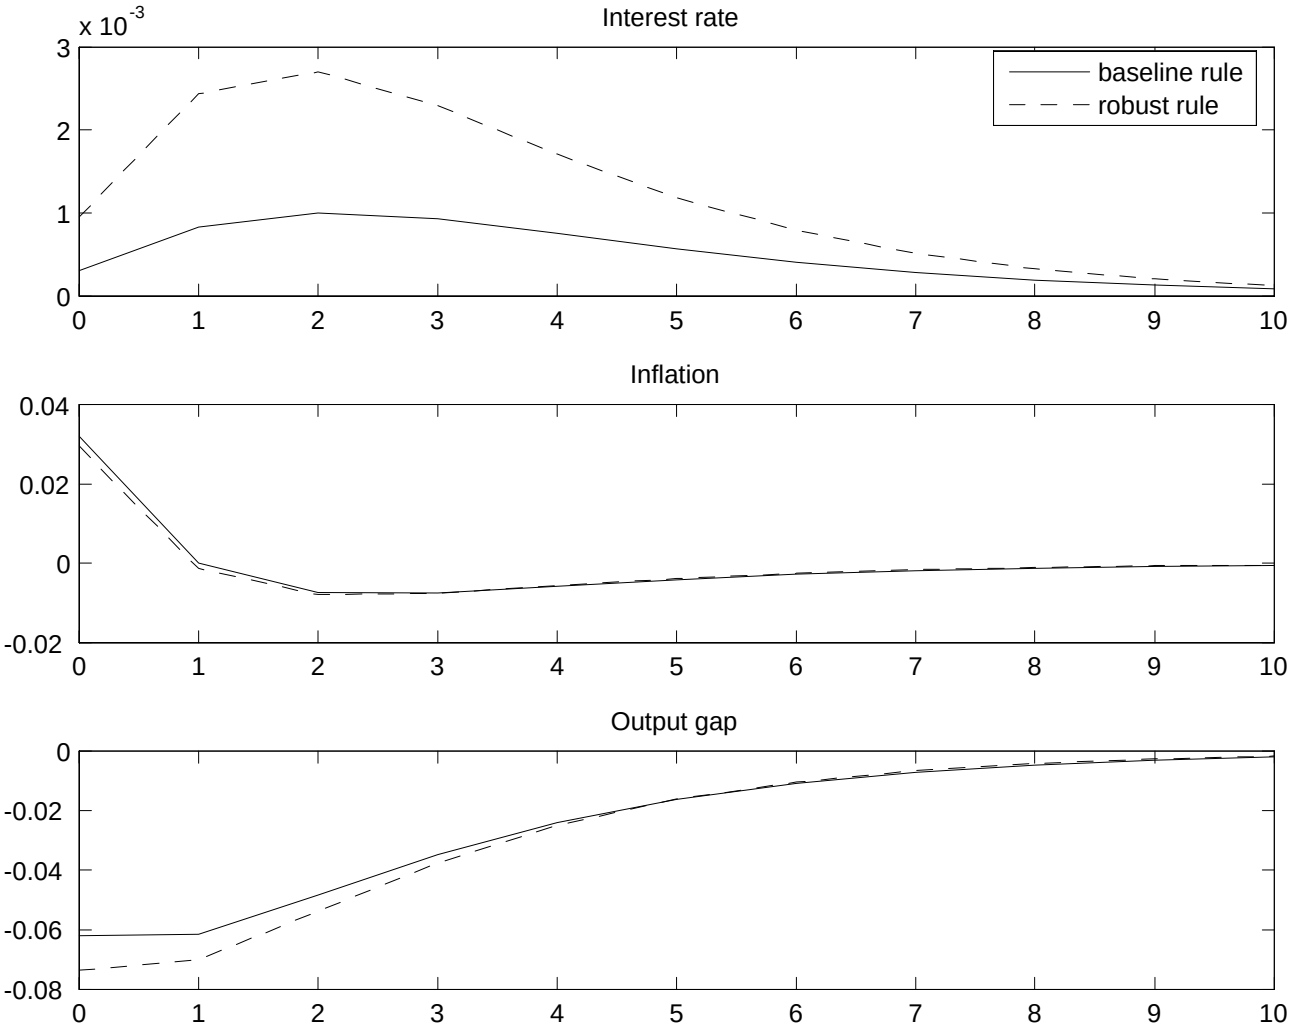

Figure 2: Impulse responses to shock μ $\left(\rho_{\mu} = 0.35\right)$.

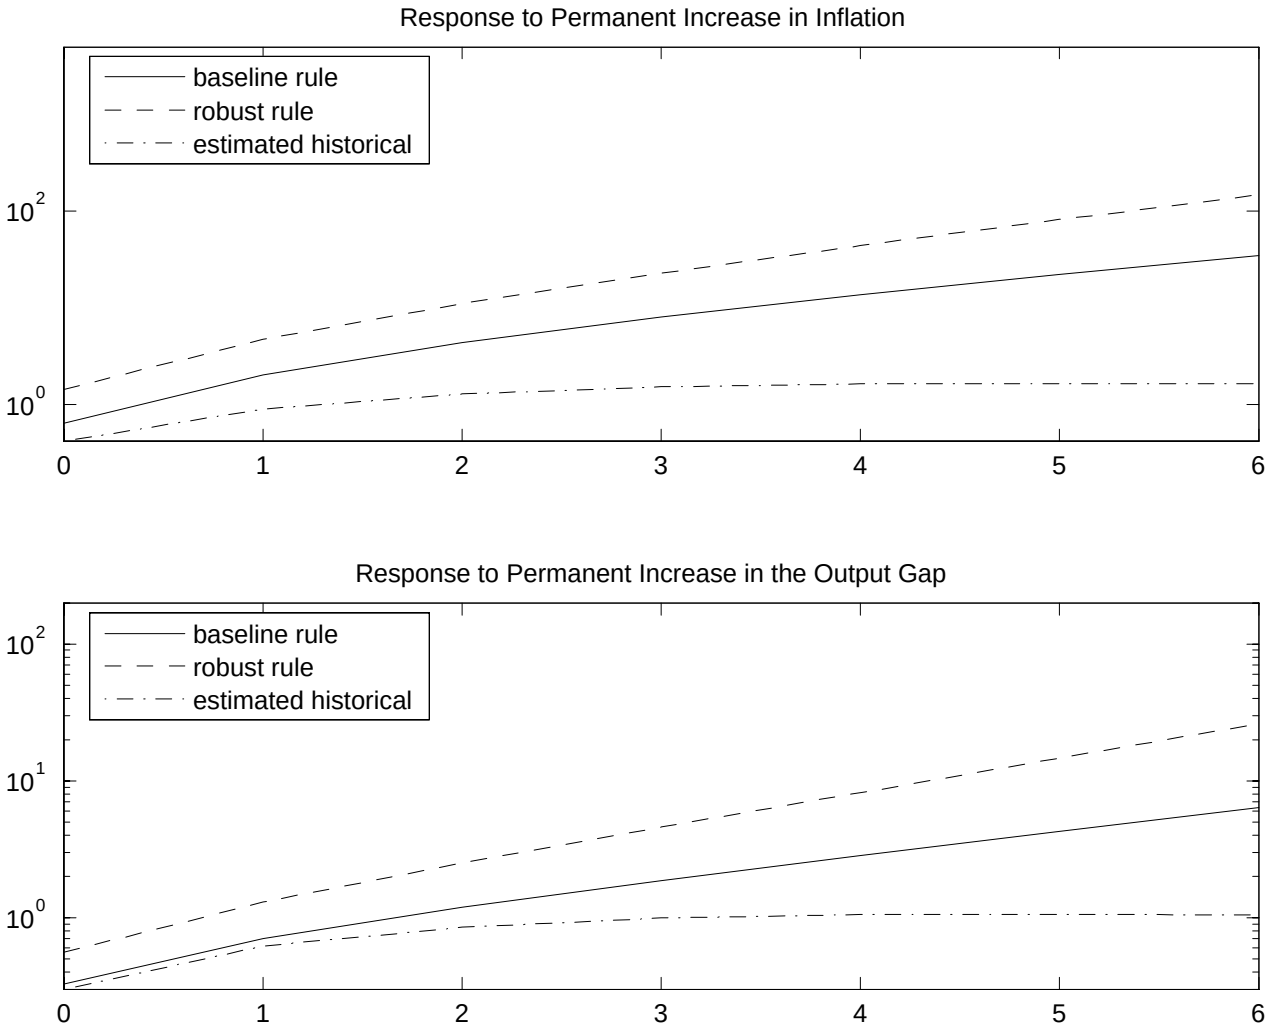

Figure 3: Interest-rate response to a permanent increase in inflation and in the output gap.

Figure 4a: δ-shock

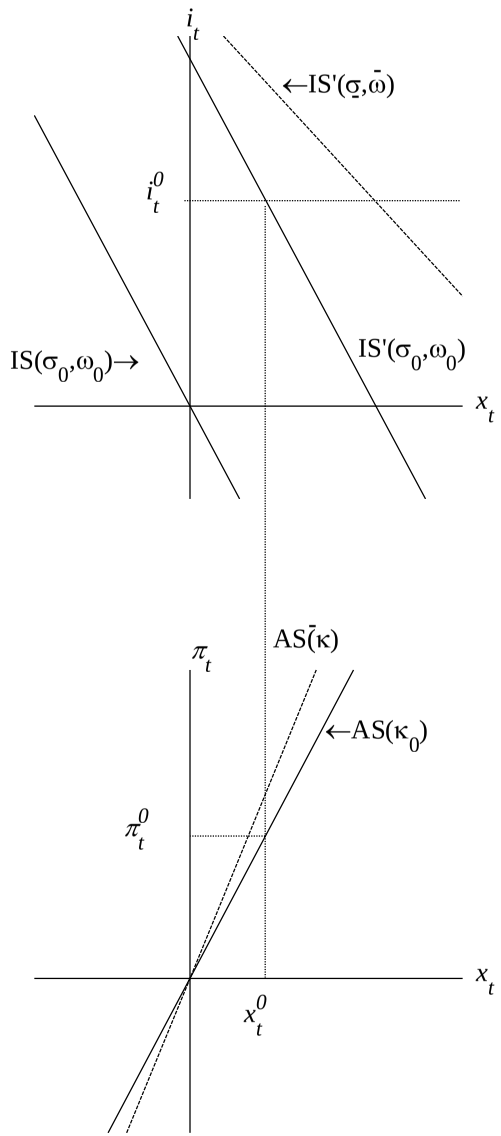

Figure 4b: ε-shock

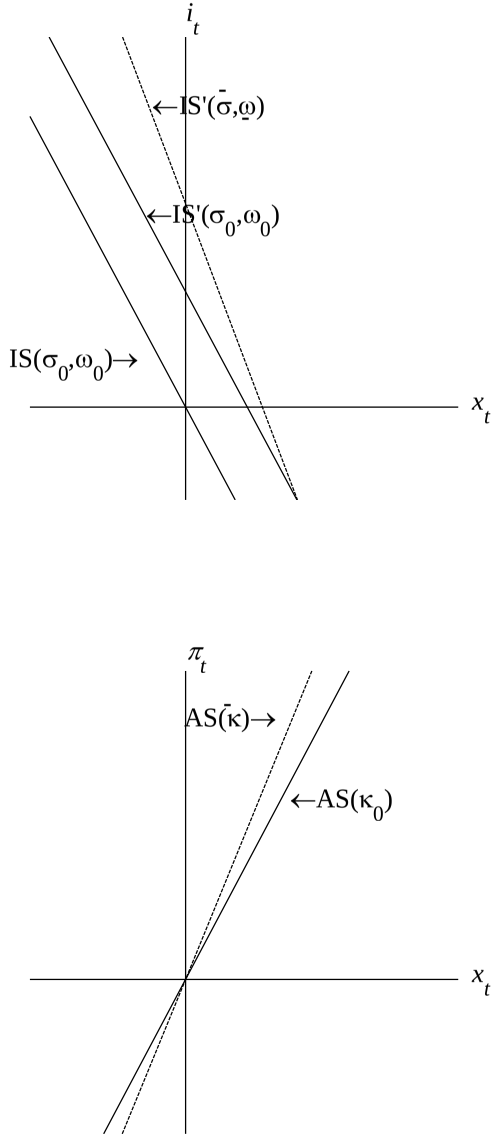

Figure 4c: μ-shock

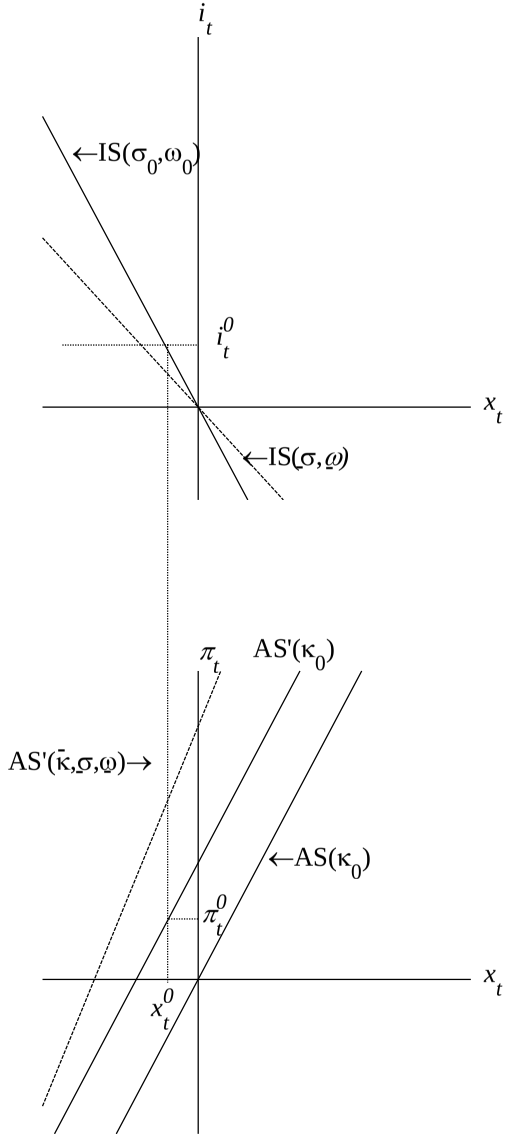

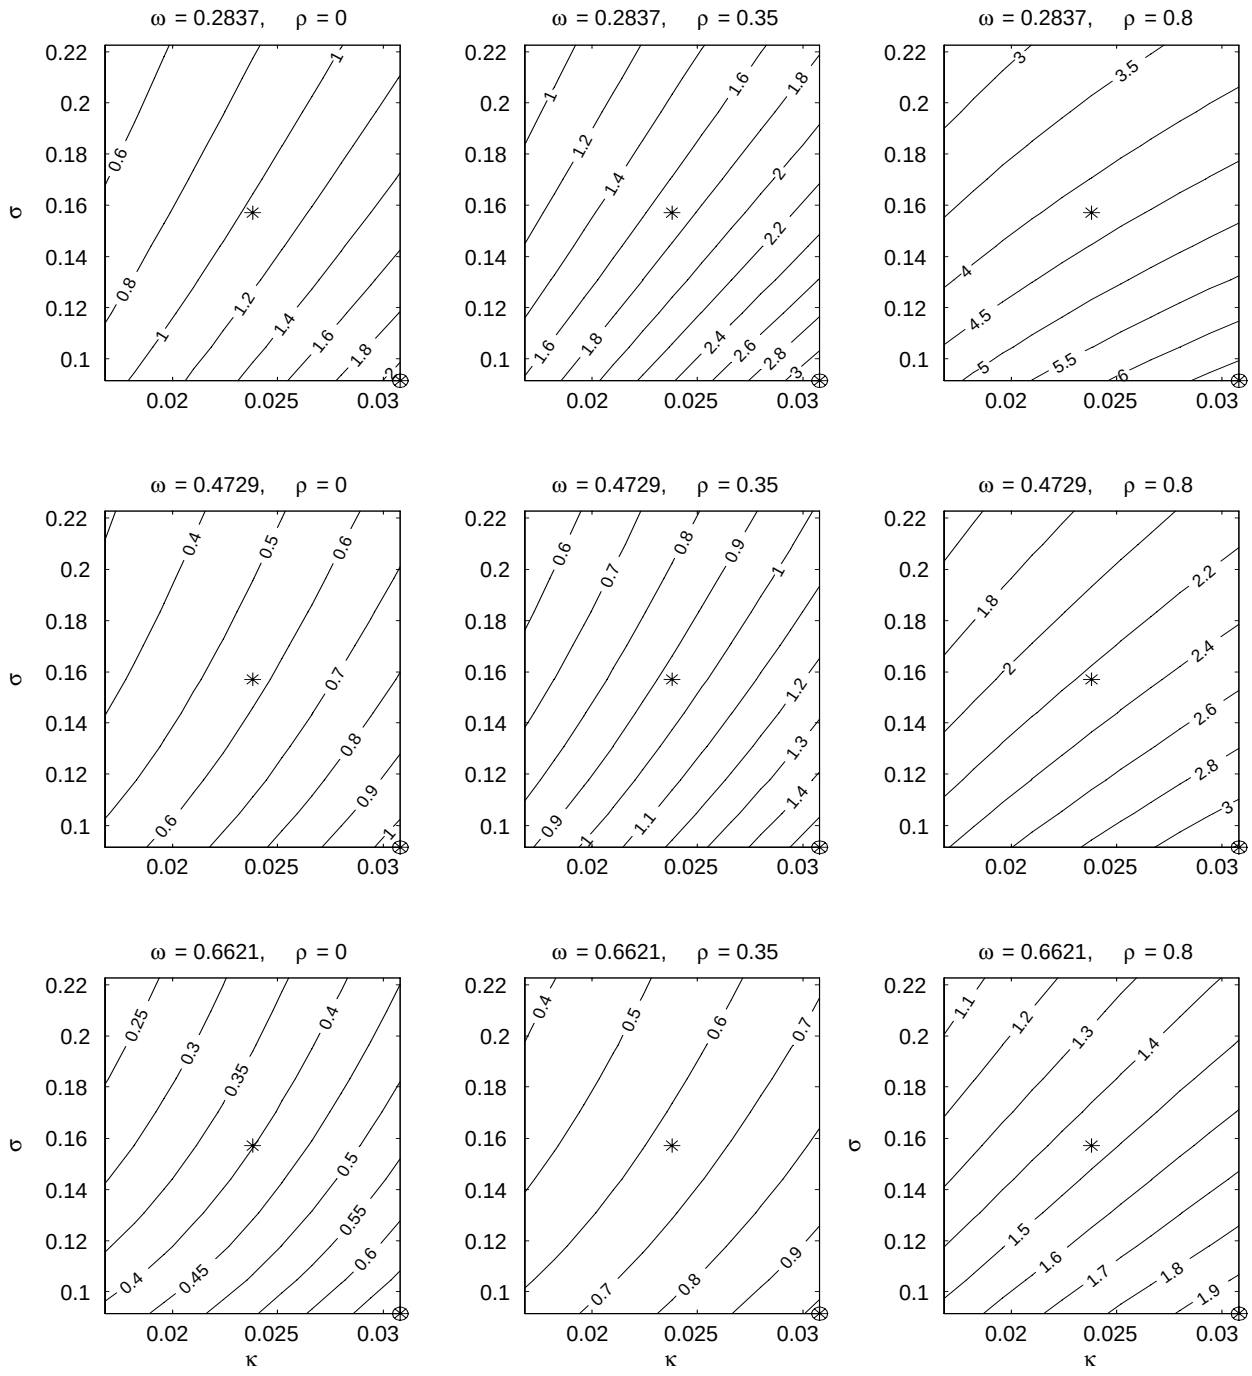

Figure 5: $E[L_0]$ with robust optimal rule in (κ, σ) plane $[v = 1]$.

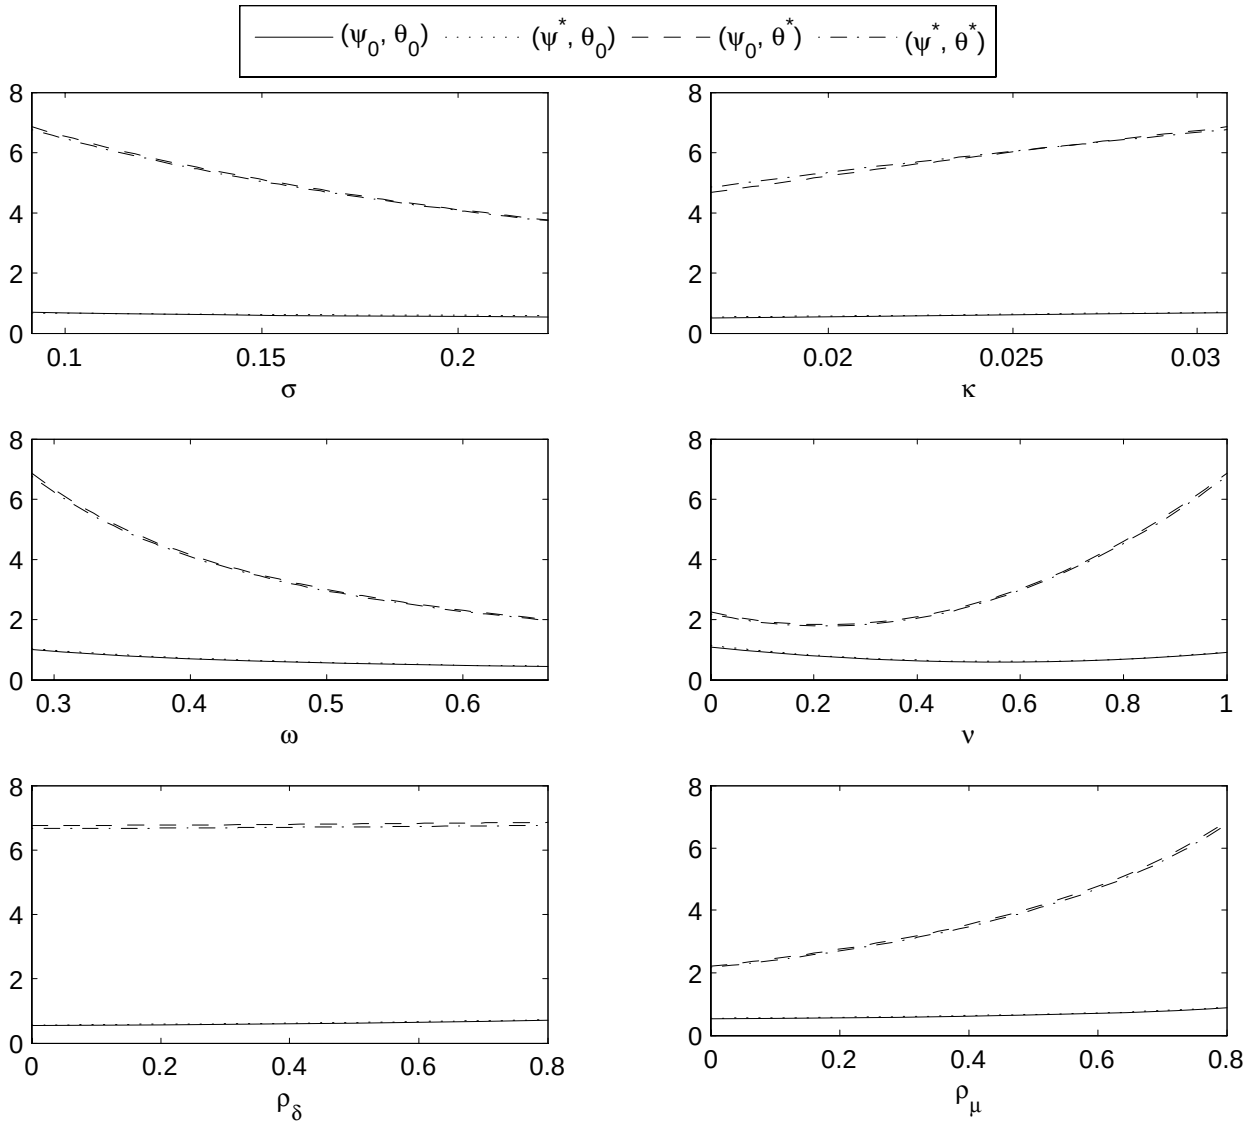

Figure 6: $E[L_0]$ with optimal $\left(\psi^0\right)$ and robust optimal (ψ^*) policy rules.

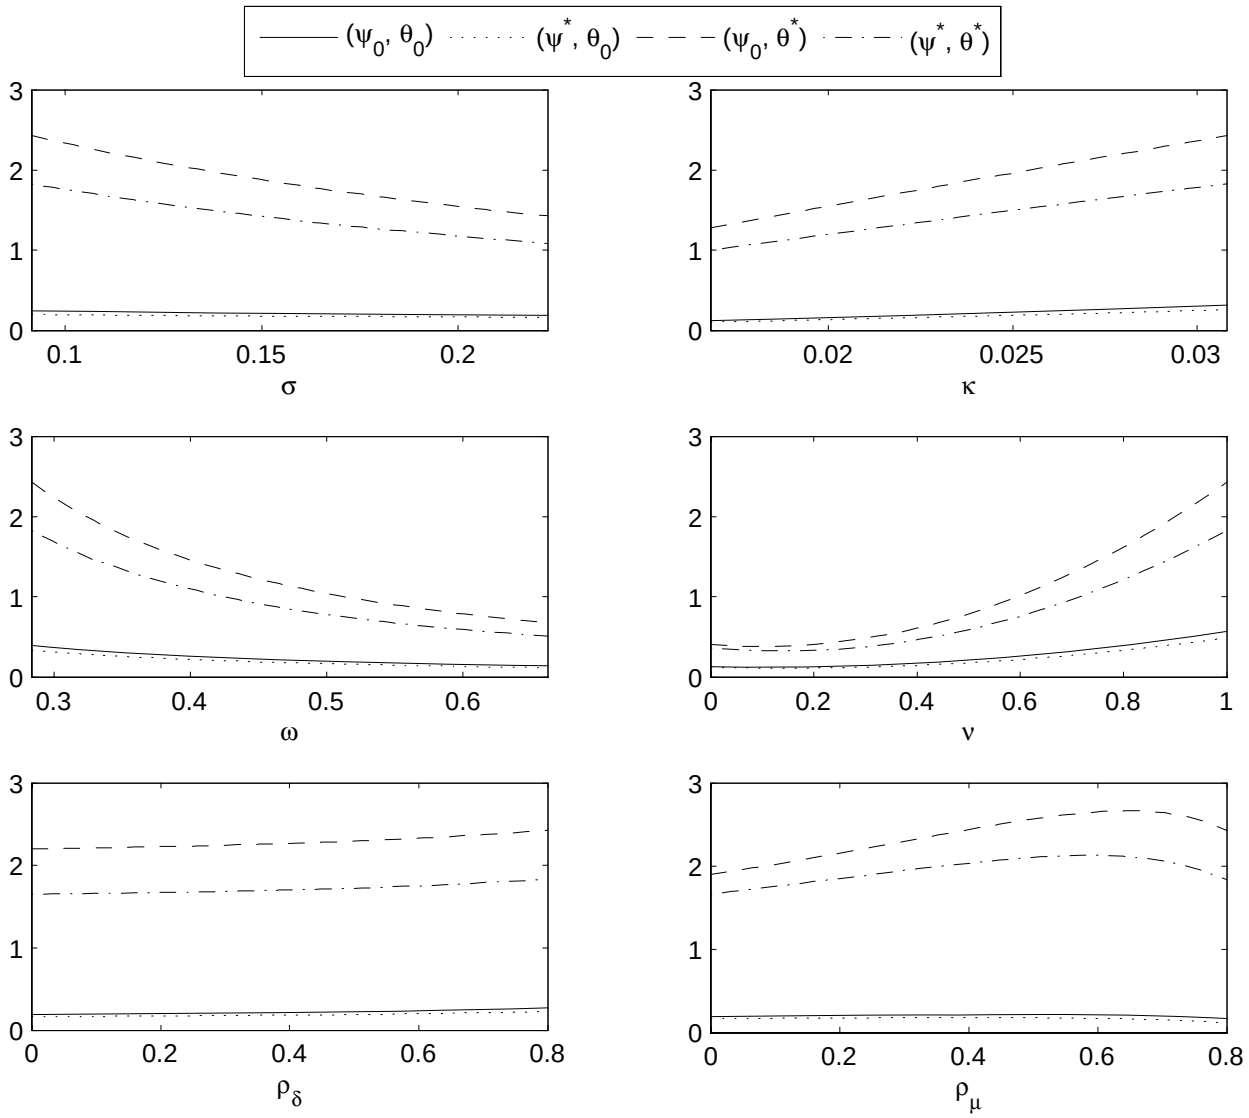

Figure 7: $V[\pi]$ with optimal $\left(\psi^0\right)$ and robust optimal (ψ^*) policy rules.

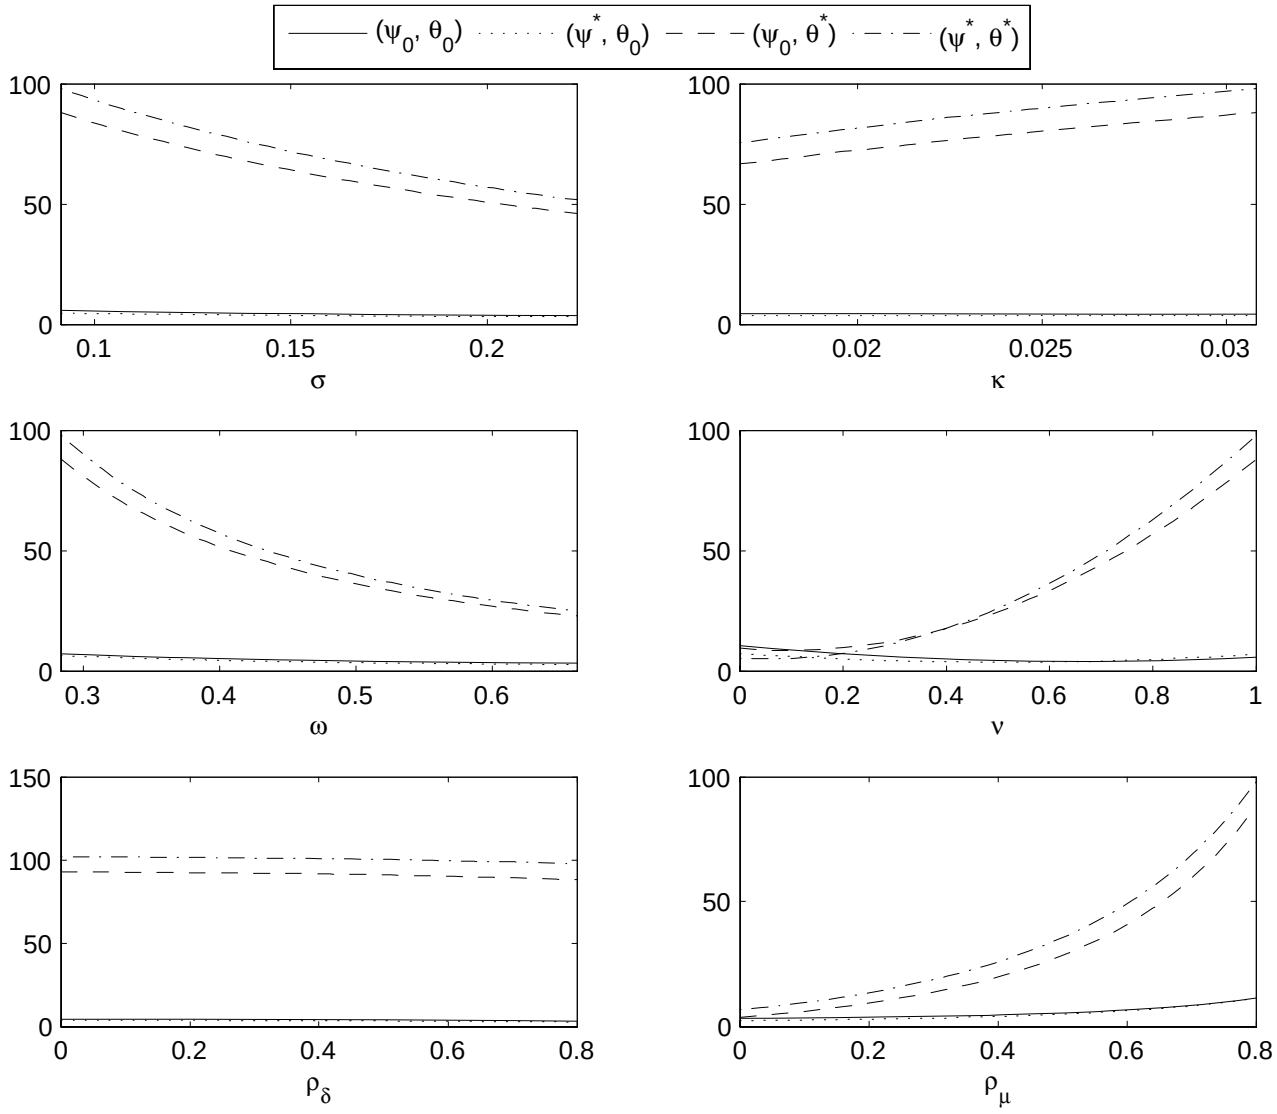

Figure 8: $V[x]$ with optimal $\left(\psi^0\right)$ and robust optimal (ψ^*) policy rules.

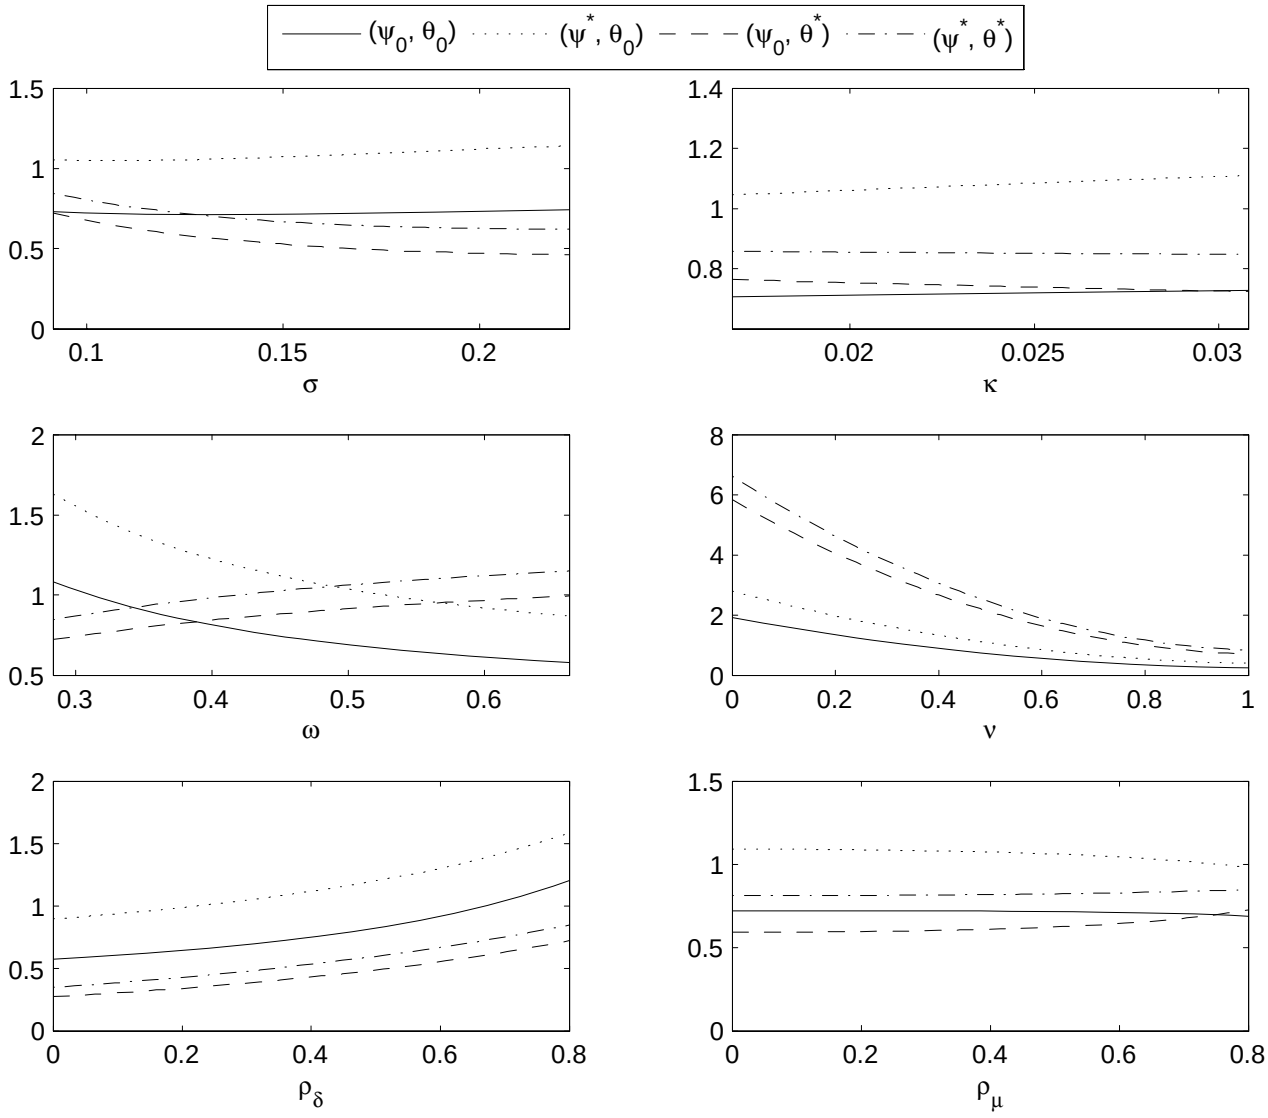

Figure 9: $V[i]$ with optimal $\left(\psi^0\right)$ and robust optimal (ψ^*) policy rules.

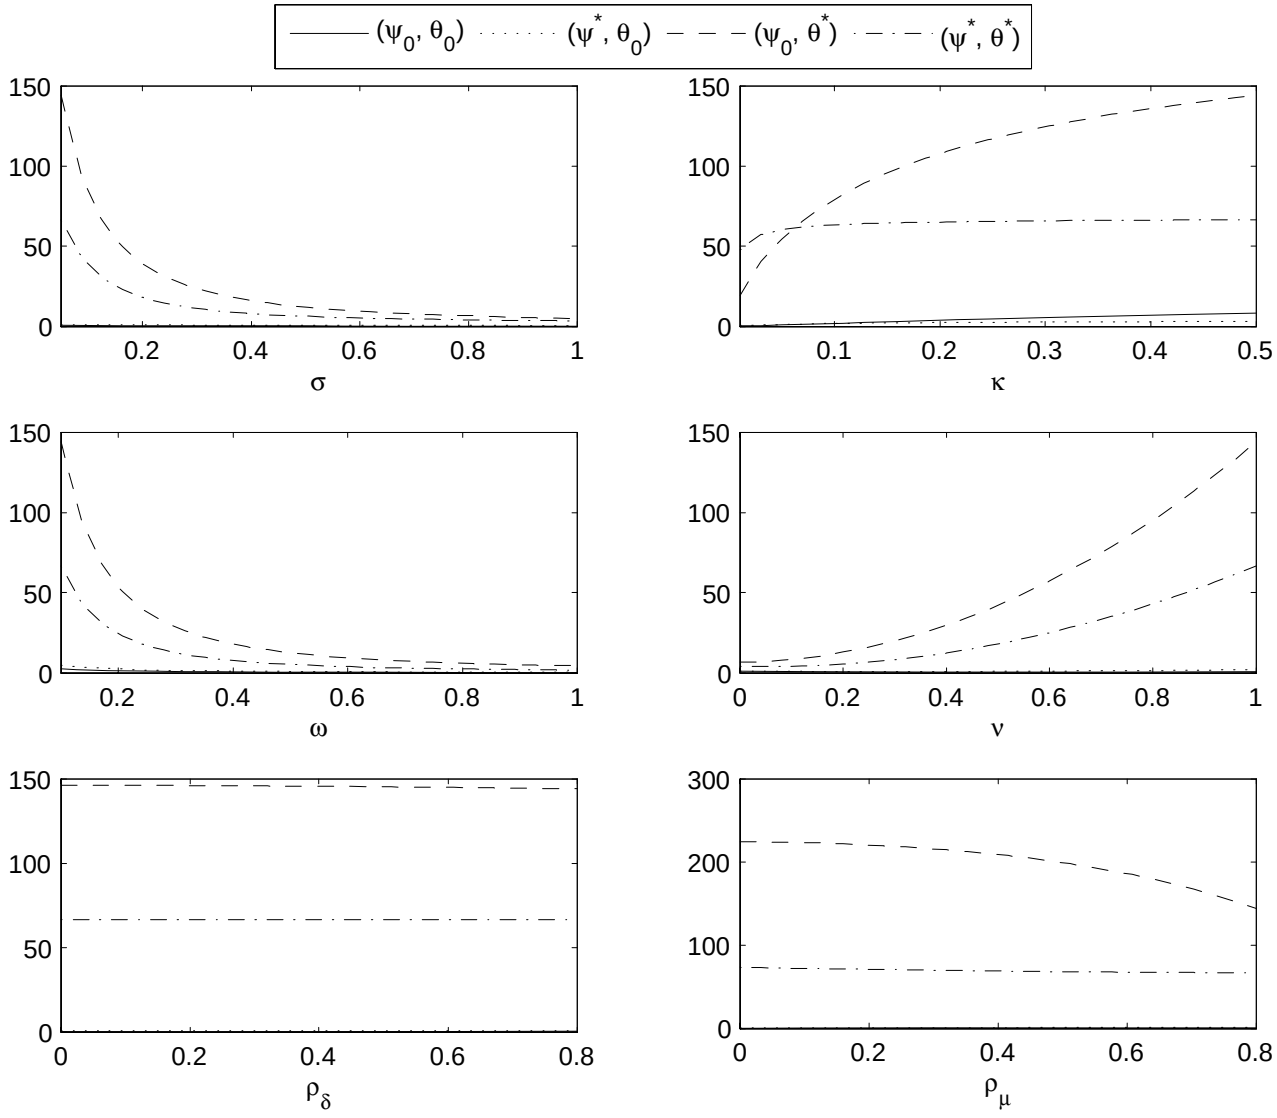

Figure 10: $E[L_0]$ with optimal $\left(\psi^0\right)$, robust optimal (ψ^*) policy rules, and large uncertainty about σ, κ, ω.