Policy evaluation and uncertainty about the effects of oil prices on economic activity

Francesca Rondina*†

University of Wisconsin - Madison

Job Market Paper

January 10th, 2009
(comments welcome)

Abstract

This paper addresses the issue of monetary policy evaluation in environments in which the policymaker is uncertain about the way oil prices affect economic performance. Despite the large literature investigating the response of economic variables to oil price shocks, there is still much debate about the predominant mechanisms by which oil prices have an impact on economic activity. In this paper, I consider models of the economy due to Solow (1980), Blanchard and Gali (2007), Kim and Loungani (1992) and Hamilton (2004), which incorporate different assumptions on the process through which oil prices are believed to affect the economy. I first study the characteristics of the model space and I analyze the likelihood of the different specifications. I show that these are equally plausible alternative models and that, as a consequence, the policymaker is faced with the problem of model uncertainty. Then, I use the Bayesian approach proposed by Brock, Durlauf and West (2003, 2007) and the minimax approach developed by Hansen and Sargent (2008) to integrate this form of uncertainty into policy evaluation. I present action dispersion and outcome dispersion analysis showing the extent to which monetary policies and their consequences are model dependent and I study the policies suggested by the minimax and minimax regret decision criteria. I find that, in the described environment, the standard Taylor (1993) rule is outperformed in terms of outcome dispersion by alternative simple rules in which the policymaker introduces persistence in the policy instrument and responds to changes in the real price of oil.

*Contacts: 7316 Sewell Social Science Building, 1180 Observatory Drive, Madison, 53706, WI-USA; phone: 001 608 263 5642, email: frondina@wisc.edu.

†I am especially grateful to Steven Durlauf for his guidance and constant encouragement. I am particularly thankful to William Brock for his comments and valuable suggestions and to Giacomo Rondina and Emmanuele Bobbio for helpful discussions and advice. I have also benefited from comments from Kenneth West, Noah Williams, Federico Diez and Nonarit Bisonyabut. All errors remain my own.

1 Introduction

This paper investigates issues related to the evaluation of monetary policy in the presence of model uncertainty. In particular, the analysis focuses on economic environments in which the policymaker is uncertain about the mechanism through which oil prices affect economic variables. In this context, this work aims to present an extensive analysis and a range of measures that can support policymakers' decision activity by providing information on the sensitivity of different policy rules to model uncertainty.

In recent years, the literature in macroeconomics has devoted large attention to the problem of model uncertainty in economic policy. In particular, this issue has received increasing interest, among economists as well as policymakers, when applied to monetary policy.[1] Some relevant contributions in this area are represented by Brock, Durlauf and West (2003, 2007), Cogley and Sargent (2005), Giannoni (2007), Hansen and Sargent (2001a, 2001b, 2008). These works develop theoretical frameworks for policy design and evaluation in uncertain environments and provide applications to different forms of uncertainty that commonly arise in monetary policy decisions.[2]

This paper applies some of the techniques developed in the literature on model uncertainty to a context in which the policymaker is uncertain about the effects of oil prices on the economy. Despite the number of contributions studying the response of economic variables to oil price shocks, there is still much debate about the mechanisms through which oil prices are believed to have an impact on economic activity. This debate originates from the fact that oil prices can indeed affect the economy in several ways. Changes in oil prices directly affect the costs of production (transportation and heating, for instance) as well as the price of goods made with petroleum products. Moreover, oil price increases are likely to increase the general price level, which can reduce employment if wages are rigid. Finally, oil price shocks can also lead to reallocation of labor and capital between sectors of the economy, and induce greater uncertainty about the future, which might reduce purchases of large-ticket consumption and investment goods. The different contributions in this area often disagree on which of these factors should be regarded as the main channel through which oil prices affect output and other economic variables.

The lack of consensus on the predominant mechanism through which oil prices affect the economy leads to different views about the ability of monetary policy to contrast the effects of oil price shocks. This generates a substantial disagreement over the way monetary policy should optimally respond to changes in oil prices.[3] In this environment, the extension of the

[1] On the policymaking side, see Dow (2004) for a description of the methodological approach that the Bank of England and the ECB have taken in response to the problem of model uncertainty.

[2] For instance, Brock, Durlauf and West (2007) present an example based on uncertainty on the way the public forms expectations on future economic variables, while in Cogley and Sargent (2005) the policymaker is uncertain about the specification of Phillips curve to be used for policy decisions.

[3] An example is provided by the recent debate between Bernanke, Gertler and Watson (1997, 2004) and

techniques developed in literature on model uncertainty seems to be quite natural, and at the same time essential to sound policymaking.

In this paper I consider the problem of a policymaker who wants to explore possible courses of action to be undertaken in response to an oil price shock. He is uncertain about the way oil prices affect the economy, and he is particularly interested in investigating the sensitivity of his policy decisions to this form of uncertainty. This paper provides an analysis of the extent to which monetary policies and their consequences are model dependent, and studies the policy recommendations of Bayesian and non-Bayesian criteria. The main finding of this paper is that, in the described environment, the standard Taylor (1993) rule is outperformed by alternative simple rules in which the policymaker introduces persistence in the policy instrument, and responds to changes in the real price of oil.

The contribution of this work to the existing literature is twofold. First, I provide an analysis of the likelihood of three alternative frameworks that have been proposed to explain the effects of oil prices on economic variables. I study optimal simple policy rules in each of these frameworks, and I provide evidence that the optimal response to a change in the price of oil is model dependent. Second, I present an application of a range of techniques developed in the model uncertainty literature to this specific form of uncertainty.

This paper is related to the literature on policy design and evaluation in uncertain environments. In recent years, two major directions of work have emerged in this area. The first one is represented by the contributions of Hansen and Sargent (2001a, 2001b, 2008). In this approach, uncertainty is defined over specifications that lie within some distance from a baseline framework, and preferences are assumed to follow a minimax rule with respect to model uncertainty.[4] A second direction is represented by the contributions of Brock, Durlauf and West (2003, 2007). In this approach, the model space includes specifications that are not close to each other according to some metric, and model uncertainty is introduced in the policymaker's decision process through the technique of Bayesian model averaging.[5] Recently, Brock, Durlauf, Nason and Rondina (2007) have proposed ways of introducing the minimax approach due to Hansen and Sargent to contexts in which the model space does not necessarily include only specifications that lie within some distance from the baseline model.

This paper is methodologically based on Brock, Durlauf and West (2003, 2007) (from now BDW, 2003, 2007) and Brock, Durlauf, Nason and Rondina (2007) (from now BDNR). The decision to follow these approaches was motivated by the fact that the uncertainty over the mechanisms through which oil prices affect economic performance is largely non-local. The

---

Hamilton and Herrera (2004) about the role of monetary policy in the economic downturns following the oil price shocks episodes of the postwar period.

[4] In more detail, the decision maker is assumed to maximize while nature minimizes over the set of models in the model space. Applications of this approach to monetary policy can be found in Giannoni (2007), Onatski and Stock (2002) and Brock and Durlauf (2004).

[5] The works of Cogley and Sargent (2005) and Cogley, Colacito and Sargent (2007) are examples of applications of this approach to the analysis of monetary policy.

description of the model space in Section 4 and 5 will provide evidence about this statement. In addition, BDW (2007) and BDNR (2007) introduce policy evaluation techniques that move beyond standard model averaging methods, and that are useful in providing a more extensive and comprehensive policy analysis. More specifically, BDW (2007) propose a range of measures and visual tools that supply the policymaker with more information than a simple summary statistic in which model dependence has been integrated out. On the other hand, BDNR (2007) introduce applications to policy evaluation of non-Bayesian approaches based on the minimax and minimax regret criteria, which have the advantage of not requiring any previous knowledge on the characteristics of the model space.

This work is also related to the large literature studying the impact of oil price changes on economic activity. The purpose of this paper is not to take a position in the debate over the different models explaining the effects of oil prices on economic performance. Rather, I show that different frameworks, based on different channels of transmission of oil price shocks, are equally plausible alternative representations of the economy. Finally, this paper is related to the literature on the response of monetary policy to changes in oil prices. Recent contributions have focused on the role of monetary policy in the downturns following the large oil price shocks of the postwar period (Bernanke, Gertler and Watson 1997, 2004; Hamilton and Herrera, 2004; Leduc and Sill, 2004) and in the milder reaction of economic variables to oil price shocks since the mid 1980s (Blanchard and Gali, 2007; Herrera and Pesavento, 2007). This work provides some additional insights in this area by explicitly showing that the consequences of monetary policy responses to changes in oil prices significantly depend on the model of the economy under consideration.

The remainder of the paper is organized as follows. Section 2 summarizes the techniques that I will use to incorporate model uncertainty into policy evaluation. Section 3 illustrates the main mechanisms that have been proposed to model the effects of oil prices in the economy. Section 4 characterizes the model uncertainty problem. Section 5 defines the model space and studies its basic properties. Section 6 reports the empirical results of the policy evaluation exercise. Section 7 concludes.

2 Incorporating model uncertainty into policy evaluation

In this section, I will summarize the approaches to policy evaluation developed in BDW (2003, 2007) and BDNR.[6] These techniques will then be applied in my policy evaluation exercise in section 6.

[6] This section only provides a brief explanation of the techniques that I will use in section 6 of the paper. For a more thorough description of these methods, see BDW (2003, 2007) and BDNR.

2.1 General Framework

The central idea in the approach proposed by BDW (2003, 2007) is that model uncertainty should be considered as a component of policy evaluation. This idea has two implications. The first one is that model uncertainty should not be resolved prior to the evaluation of a policy rule through the selection of a specific model of the economy. The second one is that policy evaluation should explicitly account for the absence of complete information concerning the specification of the true data-generating process.

Consider the problem of a policymaker who is interested in evaluating the effect of a policy rule p on an outcome θ. In this context, policies are typically evaluated based on the conditional probability measure:

$$\mu\left(\theta \mid m, p, \beta_m\right) \tag{1}$$

where m denotes a model and β_m is a vector of parameters that indexes the model. If the model m is known, the available data d can be used to estimate the vector of parameters β_m. In this case, (1) can be rewritten as:

$$\mu\left(\theta \mid m, p, d\right) \tag{2}$$

The approach to policy evaluation in uncertain environments proposed in BDW (2003, 2007) entails computing the probability measure $\mu\left(\theta \mid d, p\right)$ from (2) by treating model uncertainty as any other from of uncertainty affecting θ. This can be done by using standard application of probability arguments to eliminate the conditioning on m in (2). Let M be the space of possible data-generating processes, then we have:

$$\mu\left(\theta \mid d, p\right) = \sum_{M} \mu\left(\theta \mid m, p, d\right) \mu\left(m \mid d\right) \tag{3}$$

where $\mu\left(m \mid d\right)$ is the posterior probability of model m given data d. By Bayes' rule, this can be characterized as follows:

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

where $\mu\left(d \mid m\right)$ is the likelihood of the data given model m and $\mu\left(m\right)$ is the prior probability assigned to model m.[7]

Let now consider a policymaker that evaluates policies according to the expected losses generated by a loss function $l\left(\theta\right)$. From the previous discussion, the analysis that incorporates model uncertainty should calculate:

$$E\left(l\left(\theta\right) \mid d, p\right) = \int_{\Theta} l\left(\theta\right) \mu\left(\theta \mid p, d\right) d\theta \tag{5}$$

[7] See BDW (2007) for an interesting discussion of some interpretations of the role of model uncertainty in policy evaluation that can be inferred from this derivation.

The empirical part of this paper will involve computation of expected losses of this form, given a standard loss function that will be defined in section 4.

The model averaging approach has some attractive properties, first and foremost the fact that it allows for the assessment and comparison of policies without conditioning on a given element of the model space. However, its implementation presents several issues, mainly related to the definition of the model space M and to the specification of the prior probabilities for its elements. See BDW (2003, 2007) for a more exhaustive discussion of the implementation issues of this approach.

2.2 Outcome dispersion and action dispersion measures

In addition to the model averaging approach, BDW (2007) propose additional ways of communicating information on the effects of different policies in an environment characterized by model uncertainty. The introduction of these statistics is motivated by several considerations. First, the policymaker might be interested in aspects of the conditional density $\mu(\theta \mid m, p, d)$ that are lost in the averaging process. Second, he might be concerned about the behavior of this conditional density only in some specific models rather than others. Third, he might be interested in knowing which policies have an outcome that is relatively more stable over the different specifications in the model space.

For all these reasons, it could be useful to enrich the policy evaluation exercise by including additional measures that are able to give a broader picture of the effects of a policy under alternative assumptions. BDW (2007) introduce two measures that provide a characterization of the extent to which monetary policies and their consequences are model dependent. These measures are *outcome dispersion* and *action dispersion.*

Outcome dispersion measures the variation in loss that occurs when different models are considered, given a fixed policy rule. In other words, this measure describes how the losses associated with a specific policy rule are model dependent, thus providing information about the robustness of the selected policy rule over different models. Action dispersion, on the other hand, measures how the optimal policy differs across alternative models in a model space. A distinct optimal policy can be computed for each model, so that a range of different policies can be obtained for the models in the space. The analysis of action dispersion provides information on the sensitivity of the optimal policy rule to model choice.

2.3 Minimax and minimax regret

In addition to the outcome dispersion and action dispersion measures, I will also consider non-Bayesian approaches based on minimax and minimax regret criteria. The central idea on which these approaches are based is that the policymaker might be interested in obtaining information about policy rules that are not optimal, but that work relatively well regardless of which model is true.

The minimax approach has been largely used by Hansen and Sargent (2001a, 2001b, 2008) as the basis for robustness analysis in macroeconomics. Following BDNR, in my policy evaluation exercise I will define the minimax policy choice as:

$$\min_{p \in P} \max_{m \in M} E\left(l\left(\theta\right) \mid p, d, m\right) \tag{6}$$

One of the main issues of the minimax approach is that of being extremely conservative, since it always assumes the worst possible scenario in assessing the different policies. The minimax regret approach has been proposed to avoid this problem. Indeed, minimax regret is based on the relative loss associated with a given policy, under the assumption that the policymaker does not know the correct model of the economy. Following again BDNR, I will define the minimax regret policy rule as:

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

where $R\left(p, d, m\right)$ is the regret function defined as:

$$R\left(p, d, m\right) = E\left(l\left(\theta\right) \mid p, d, m\right) - \min_{p \in P} E\left(l\left(\theta\right) \mid p, d, m\right) \tag{8}$$

Given a model, the regret function measures the loss suffered by a policy relative to the loss under the optimal policy for that specific model. The definition of the regret function illustrates how this criterion is able to avoid the problems associated with models that comport relatively high losses, regardless of the choice of the policy rule.

BDNR offer a more comprehensive description of the properties of the minimax and minimax regret criteria and provide examples of applications of these approaches that have been proposed in the literature.

3 Modeling the effects of oil prices on the economy

This section provides a brief review of the most relevant contributions on the effects of oil prices on economic activity.[8]

The literature in economics has proposed many different mechanisms through which oil prices can affect economic performance. Some early studies, such as Solow (1980) and Pindyck (1980) focused on the demand-side effects of changes in oil prices. In these frameworks, the direct and immediate effect of a change in oil prices is a change in the overall price level, which in turn has an effect on employment and other real variables due to the Keynesian

[8] Extensive reviews of the different mechanisms by which oil prices can affect economic performance are provided by Mork (1994) and Segal (2007). See also Hamilton (2002).

assumption of rigid wages. It follows that, in these models, the main channel through which oil price variations affect output is wage rigidities. A similar approach is the one proposed by Blanchard and Gali (2007), in which the assumption of price rigidities is added to that of wage rigidities.

A second strand of literature considers the supply-side effects of changes in oil prices. These works are usually based on a production function in which energy is one of the inputs, so that an exogenous change in the price of oil affects output directly by changing productivity and employment through a change in the wage level. Some contributions based on this mechanism are Rasche and Tatom (1977) and Kim and Loungani (1992). This way of explaining the effect of oil prices on output seems to be quite natural in the context of a standard neoclassical economic model. Several other papers have considered departures from the standard neoclassical framework that are able to explain additional indirect effects of an oil price shock on output. For instance, Finn (2000) focuses on the impact of changing capacity utilization rates, while Rotemberg and Woodford (1996) consider a model characterized by imperfect competition, in which additional effects on output are caused by changes in business markups.

Finally, one last group of contributions has focused on the effects of oil price shocks on short-run economic performance as the consequence of allocative disturbances. Some examples of this literature are Bernanke (1983) and Hamilton (1988). These studies have the relevant feature of suggesting a nonlinear relation between oil prices and output. An oil price increase will decrease demand for some goods but possibly increase demand for others. As a consequence, if it is costly to reallocate labor or capital between sectors, then an oil shock will be contractionary in the short run. However, an oil price decrease would require the same type of reallocative process, and for this reason it could be contractionary as well in the short run.

4 Model Uncertainty

I consider the problem of a policymaker who wants to investigate possible policy responses to oil price changes. He knows that many different mechanisms have been proposed in the economic literature to explain the effects of oil prices on economic activity. In particular, he believes that the true model of the economy might be one of the following three frameworks:

- Solow (1980) (from now on denoted as S), in which the most significant effect of a change in oil prices is a change in the overall price level, which in turn affects employment and real variables due to the assumption of nominal wage rigidities. Therefore, in this model the main channel through which oil prices have an impact on output is nominal wage rigidities.

- Blanchard and Gali (2007) (from now on denoted as BG), in which the central effect of a change in oil prices is a change in the overall price level, which in turn affects employment and real variables due to the assumption of price and real wage rigidities. This is a new Keynesian type of model, and price rigidities are introduced in the economy by assuming Calvo pricing. In this framework, the channel through which oil prices have an impact on economic activity is real wage and price rigidities.

- Kim and Loungani (1992) and Hamilton (2005) (from now on denoted as H), in which changes in the price of oil affect output directly by changing productivity and have an impact on employment through a change in the wage level. This is a standard neoclassical type of model, characterized by perfect competition and flexible prices and wages.

Given his beliefs on the possible true models of the economy, the policymaker considers three different econometric frameworks that incorporate the main features of each one of these models. Each framework consists of two equations, one for the output gap and one for the inflation rate, and includes the following variables: y_t, which represents output gap; π_t which is core CPI inflation, i_t which is the interest rate (the policy instrument), and s_t which is the real price of oil.[9]

The S model is represented by the following equations:

$$y_t = \alpha_y^S (L) y_{t-1} + \alpha_\pi^S [\pi_{t-1} - E_{t-2}(\pi_{t-1})] + \alpha_s^S (L) s_{t-1} + \omega_{y,t}^S \tag{9}$$

$$\pi_t = \beta_\pi^S (L) \pi_{t-1} + \beta_y^S (L) y_{t-1} + \beta_r^S \bar{i}_{t-1} + \beta_s^S (L) s_{t-1} + \omega_{\pi,t}^S \tag{10}$$

where the effects of oil prices on output through nominal wage rigidities are captured by the unanticipated inflation term in the output equation. This econometric model can be interpreted as an example of a setup in which nominal wages are set in advance, as for instance in Woodford (2003).

The BG model is represented by the following equations:

$$y_t = \alpha_y^{BG} (L) y_{t-1} + \alpha_r^{BG} \left[\bar{i}_{t-1} - E_{t-1}(\pi_t)\right] + \alpha_s^{BG} (L) s_{t-1} + \omega_{y,t}^{BG} \tag{11}$$

$$\pi_t = \beta_\pi^{BG} (L) \pi_{t-1} + \beta_y^{BG} (L) y_{t-1} + \beta_s^{BG} (L) s_{t-1} + \omega_{\pi,t}^{BG} \tag{12}$$

which represents an example of a new Keynesian type of framework, and is similar to the model used in Rudebusch and Svensson (1999) under the assumption of backward expectations, with the only difference being the addition of the real price of oil in both equations.

[9] The use of core CPI inflation follows Blanchard and Gali (2007). See Appendix 1 for a more detailed description of the data used in the empirical analysis.

Finally, the H model is represented by the following equations:

$$y_t = \alpha_y^H(L)\, y_{t-1} + \alpha_s^H(L)\, s_{t-1} + \omega_{y,t}^H \tag{13}$$

$$\pi_t = \beta_y^H(L)\, y_{t-1} + \beta_\pi^H(L)\, \pi_{t-1} + \beta_r^H \bar{i}_{t-1} + \beta_s^H(L)\, s_{t-1} + \omega_{\pi,t}^H \tag{14}$$

which are characterized by the independence of real variables from money and inflation. This econometric model represents an example of a Sidrauski-Brock type of model, or of a model with perfect competition and complete financial markets (see again Woodford, 2003). Equation (13) has been frequently used by Hamilton (2003, 2005) to estimate the effects of oil prices on output.

In all specifications the policy instrument i_t is assumed to affect the economy in the form of the average annual rate: $\bar{i}_t = \frac{1}{4}\sum_{j=1}^{4} i_{t-j}$. In the BG model, this assumption has been used in the literature on the new Keynesian Phillips curve (see, for instance, Rudebusch and Svensson, 1999). In the S and H models, I decided to use $\bar{i}_t$ as well, in order to have some consistency in the variable through which monetary policy affects the economy.

Each model is completed with the definition of a policy rule for the interest rate i_t and with the specification of a process for the real price of oil s_t. These are common to all frameworks and are defined next.

4.1 The policy rule

In regard to the policy rule, I assume that the policymaker employs a simple nominal interest rate rule in the form:

$$i_t = g_\pi \pi_t + g_y y_t + g_i i_{t-1} + g_s s_t \tag{15}$$

This rule is similar to the one used in BDW (2007) and in many other contributions in the literature on monetary policy, with the only difference being the addition of the last term, which represents the response of the policymaker to changes in the real price of oil.

Following standard assumptions in the monetary rules literature, the policymaker chooses the parameters g_π, g_y, g_i and g_s in (15) to minimize the expected loss function:

$$R = var\,(\pi_\infty) + \lambda_y var\,(y_\infty) + \lambda_i var\,(\Delta i_\infty) \tag{16}$$

In the policy evaluation exercise in section 6, different values of R will be calculated based on alternative conditioning assumptions made via specification of a policy and/or a model. I will assume that $\lambda_y = 1$ and $\lambda_i = 0.1$ as in BDW (2007); this choice is consistent with the literature using similar loss functions, see for instance Levin and Williams (2003).

4.2 The process for the real price of oil

I assume that the real price of oil follows the exogenous AR(1) process:

$$s_t = \delta_t + \rho s_{t-1} + o_t \tag{17}$$

where the intercept δ_t is allowed to change over time, so that the real price of oil takes the form of a mean changing process. This representation for the real price of oil aims to capture the nonlinearities that seem to characterize the behavior of the real price of oil.[10] In this framework, o_t represents the oil price shock, which is assumed to be uncorrelated over time, and to have mean zero and constant variance σ_o^2. The process described in (17) is a generalization of the representation used in Blanchard and Gali (2007), which simply set $\delta_t = 0$ at any time t.

The process in (17) can be rewritten in matrix form as:

$$s_t = d_t' \Phi_t + o_t \tag{18}$$

where $d_t = [\delta_t \quad \rho]'$ and $\Phi_t = [1 \quad s_{t-1}]'$. I will model changes in the vector of coefficients d_t by assuming that:

$$d_t = d_{t-1} + \Lambda_t \tag{19}$$

where Λ_t is an i.i.d. Gaussian random vector with mean zero and covariance matrix V, which is assumed to be uncorrelated with the oil price shock o_t. Equation (17) implies that the intercept δ_t changes over time while the slope ρ does not. As a consequence, Λ_t will take the form:

$$\Lambda_t = \begin{bmatrix} \varepsilon_t & 0 \\ 0 & 0 \end{bmatrix} \tag{20}$$

The policymaker believes that the true process for the real price of oil drifts over time; for this reason, he will continuously adapt its parameter estimates with nonvanishing weight on new observations. The details on the procedure used to estimate (17) are provided in Appendix 2.

5 The model space

In his policy evaluation exercise, the policymaker wants to consider different forms of model uncertainty.

- **Theory uncertainty.** The first, and most important, form of model uncertainty the

[10] See Pindyck (1998) for a discussion. Blanchard and Gali (2007) also suggest that the real price of oil would be better described by a nonstationary process.

policymaker is concerned about is theory uncertainty. Theory uncertainty refers to the policymaker's imperfect knowledge of the mechanism through which oil prices affect economic activity. This form of uncertainty is represented by the three different frameworks described in the previous section. These frameworks will originate three different classes of models that will define the model space: $M = \left\{M^{S}, M^{BG}, M^{H}\right\}$.

- **Specification uncertainty**. For each one of the three frameworks described in the previous section, the policymaker is also uncertain about the way the model should be specified. This form of uncertainty reflects the imperfect knowledge about the correct specification of the econometric framework to be estimated, which would affect the policymaker's decision process even if he knew the true model of the economy.

  In more detail, I assume that specification uncertainty refers to the number of lags of the variables of interest to be included in the estimation of each of the different models. The policymaker will incorporate this form of uncertainty by estimating equations (9), (11) and (13) with one, two, three and four lags of y_t and s_t. In the same way, he will estimate equations (10), (12) and (14) with one, two, three and four lags of y_t, π_t and with zero, one, two, three and four lags of s_t. This asymmetry in the treatment of the variable s_t arises from the possibility that core CPI inflation is not affected by changes in the real price of oil, which is a possibility that is considered at least in the BG model.[11]

5.1 Basic properties of the model space

The different forms of uncertainty that the policymaker decided to consider in his policy analysis result in a model space composed of 3,840 models, 1,280 for each group of models $M^j; j = S, BG, H$.

In estimating the different specifications, I assumed that the public forms expectations using a backward-looking approach, so that: $E_{t-1}(\pi_t) = \frac{1}{4}\sum_{j=1}^{4}\pi_{t-j}$.[12] In addition, in all specifications of (12), I assumed that $\sum_{j=1}^{J}\beta_{\pi j}^{BG} = 1$, where J is the total number of lags in the polynomial $\beta_{\pi}^{BG}(L)$. This assumption is consistent with the theory on a vertical long run Phillips curve.

The process for the real price of oil in $(17) - (20)$ was estimated using the Kalman filter learning algorithm described in Appendix 2. The choice of this learning rule, rather than the more commonly used recursive least squares rule was motivated by the fact that the Kalman

[11] Solow (1980) also acknowledges the possibility that oil price shocks do not affect core inflation, even if he says that this situation is very unlikely.

[12] The analysis of the forward looking expectations case is one of the extension I would like to investigate in the future. Another interesting extension would be to consider uncertainty on the way expectations are formed, as in BDW (2007).

filter learning algorithm discounts past data more rapidly, thus allowing the policymaker to update his beliefs on the coefficient δ_t more quickly.[13] The policymaker is ultimately interested in the estimated value of ρ and in the values of σ_o^2 and V, since these are the variables that will be used to compute the expected losses in the policy evaluation exercise. The value for ρ obtained from estimating the mean changing process in (17) is 0.9561, which is lower than the value proposed in Blanchard and Gali (2007) (0.97).

I estimated all the models in M using ordinary least squares. The data used in the estimation is described in Appendix 1. Then, I used the estimated models to obtain posterior model probabilities, which were computed using (4), and assuming an uniform prior for all the models in M, so that $\mu(m) = 1/3840$ for each model. The use of priors that do not put more weight on some models rather than others reflects the assumption that the policymaker has no previous knowledge on the true model of the economy, and thus decides to assign an equal prior probability to all of them.

Figure 1. Posterior probabilities

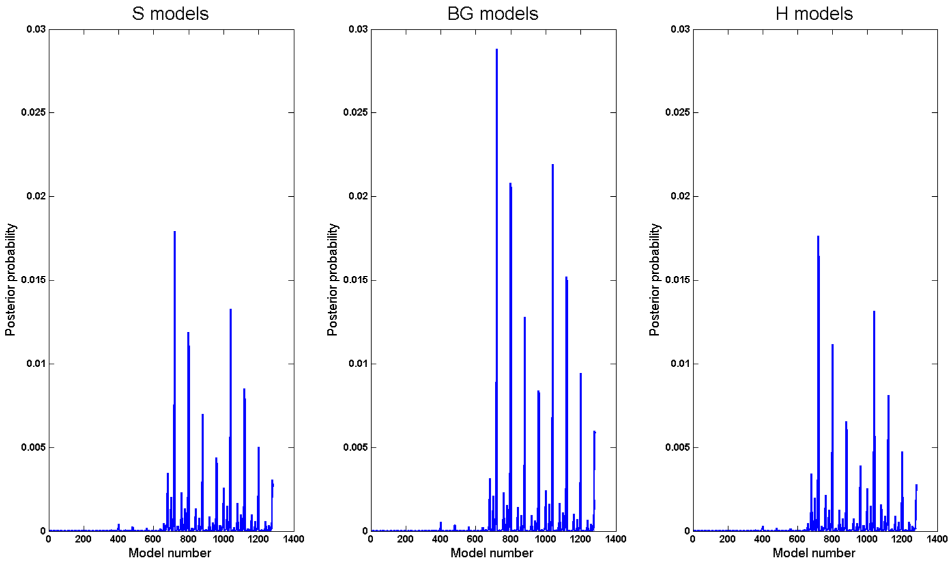

Notes: 1. Posterior probabilities for each model in the model space M. Each panel represents a class of models. As explained in the main text, these probabilities correspond to model specific BIC-adjusted likelihoods.

2. The sum of posterior probabilities over the three panels is equal to one. The sum of posterior probabilities in each panel (i.e. for any class of models) is reported in Table 1.

3. Model numbers are explained in Appendix 1.

[13] In addition, Sargent and Williams (2005) show that these two learning rules are closely related, since the recursive least squares learning rule can be approximated by a Kalman filter rule when V is set proportional to $\sigma_o^2 E(\Phi\Phi')^{-1}$.

Figure 1 reports the posterior probabilities for each model in the model space. Posterior model probabilities have been computed using the approximation suggested in Raftery (1995), so they represent BIC adjusted likelihoods. This picture shows that each of the theories described in the previous section is an equally plausible alternative representation of the economy. This picture also shows that a policymaker who decided to engage in a model selection exercise using model specific likelihoods, would find it quite difficult to discard any of these theories, since they all include specifications with BIC adjusted likelihood that is significantly different from zero.

Table 1 provides additional information on the sum of posterior probabilities for each class of models in the model space.

Table 1 - Sum of posteriors for each class of models

	M^S	M^{BG}	M^H
Sum of posterior probabilities	0.297	0.420	0.283

Table 2 reports the estimated coefficients for the specification with the highest posterior probability in each class of models. As I mentioned before, posterior probabilities are model specific BIC adjusted likelihoods. It follows that the specifications represented in Table 2 correspond to those that would have been selected within each class using BIC as the selection criterion. In each class of models, the specification with the highest posterior probability includes the following number of lags: 3 lags of the output gap, and only 1 lag of the real price of oil in the output equation; 4 lags of the output gap, 4 lags of inflation and 3 lags of the real price of oil in the inflation equation.

The specifications reported in Table 2 are those that I will use in the second part of this section to compute the simple policy rules that will be employed in the policy evaluation exercise in section 6.

Table 3 provides some summary statistics on the distribution of posterior probabilities across the model classes. Following BDW (2007), I compute the relative likelihood of a model within a class, defined as:

$$P_m = \frac{\widehat{L}_m}{\sum\limits_{m \in C} \widehat{L}_m} \tag{21}$$

where $\widehat{L}_m$ is the BIC-adjusted likelihood for model m, and C is equal to M^S, M^{BH} or M^H depending on the class under consideration.

Table 2 - Parameter estimates for models with highest posterior probabilities

(*A*) Output equation

	α_{y_1}	α_{y2}	α_{y3}	α_r	α_π	α_{s_1}	$\overline{R}^2$	DW	$s.e.$
S models	1.085 (0.005)	−0.031 (0.011)	−0.203 (0.005)	*n.a.*	0.038 (0.001)	−0.005 (0.0000)	0.90	1.95	0.57
BG models	1.092 (0.005)	−0.026 (0.011)	−0.208 (0.005)	−0.032 (0.001)	*n.a.*	−0.004 (0.0000)	0.90	1.97	0.57
H models	1.098 (0.005)	−0.028 (0.011)	−0.210 (0.005)	*n.a.*	*n.a.*	−0.005 (0.0000)	0.90	1.98	0.57

(*B*) Inflation equation

	β_{y_1}	β_{y2}	β_{y3}	β_{y4}	β_{π_1}	β_{π_2}	β_{π_3}	$\beta_{\pi 4}$	β_r	β_{s_1}	β_{s_2}	β_{s_3}	$\overline{R}^2$	DW	$s.e.$
S models	0.467 (0.023)	−0.293 (0.049)	0.323 (0.046)	−0.323 (0.020)	0.170 (0.005)	0.217 (0.005)	0.407 (0.005)	0.192 (0.006)	−0.043 (0.004)	0.010 (0.0001)	0.010 (0.0002)	−0.019 (0.0001)	0.70	2.03	2.27
BG models	0.492 (0.023)	−0.303 (0.049)	0.314 (0.045)	−0.345 (0.020)	0.175 (0.005)	0.219 (0.005)	0.413 (0.005)	0.193 (0.005)	*n.a.*	0.010 (0.0001)	0.010 (0.0002)	−0.021 (0.0001)	0.72	2.04	2.26
H models	0.467 (0.023)	−0.293 (0.049)	0.323 (0.046)	−0.323 (0.020)	0.170 (0.005)	0.217 (0.005)	0.407 (0.005)	0.192 (0.006)	−0.043 (0.004)	0.010 (0.0001)	0.010 (0.0002)	−0.019 (0.0001)	0.70	2.03	2.27

Notes: 1. Panel (A) presents the estimated coefficients for equations (9), (11) and (13) and Panel (B) the estimated coefficients for equations (10), (12) and (14) for the specification with the highest posterior probability in each class of models. Constant terms were included in all the regressions.

2. In Panel (A), output gap is the dependent variable, α_{yj} is the coefficient on output gap at lag j, α_r and α_π are the coefficients on the annual real interest rate and unanticipated inflation respectively, and α_{sj} is the coefficient on the real price of oil at lag j. In Panel (B), inflation is the dependent variable, β_{yj} is the coefficient on output gap at lag j, $\beta_{\pi j}$ is the coefficient on inflation at lag j, β_r is the coefficients on the annual real interest rate, and β_{sj} is the coefficient on the real price of oil at lag j.

3. The sample is composed of quarterly data from 1960:I to 2008:II, for a total of 194 observations. Inflation is the annualized change in core CPI; the output gap is computed using real GDP and the CBO estimate of potential GDP; the interest rate is the average Federal funds rate; the real price of oil is the difference between the log of the nominal price of oil and the log of core CPI. Additional information on the data used in the estimations is provided in Appendix 1.

The measure computed using (21) allows the policymaker to identify the models that have the highest relative posterior probability in each class. These models will be defined as those for which P_m is at least 1/50 of the model with the highest P_m within the class. In the reminder of the paper, I will denote this group of models as those having "high" likelihood or "high" posterior probability.[14]

Table 3 - Relative likelihood P

	S	BG	H
(1) Minimum P	3×10^{-22}	4×10^{-31}	2×10^{-22}
(2) Q1 P	4×10^{-16}	3×10^{-24}	2×10^{-16}
(3) Median P	2×10^{-10}	3×10^{-12}	2×10^{-10}
(4) Q3 P	6×10^{-5}	3×10^{-5}	5×10^{-5}
(5) Maximum P	0.060	0.069	0.062
(6) No. models with $P > (max\ P)/50$	108	81	105
(7) Sum of P for models with $P > (max\ P)/50$	0.9102	0.9206	0.9107
(8) Sum of P for models in top quartile	0.9970	0.9990	0.9975
(9) Sum of P	1	1	1

Note: The relative likelihood P is defined by (21). The sum of P for each class of models equals one by construction.

5.2 Optimal simple rules

I will now turn to the study the optimal simple rule in the form of (15) for the specification with the highest posterior probability in each class of models. These specification are those described in Table 2.

The purpose of this exercise is twofold. First, I want to investigate whether the three classes of models described in section 4 imply different optimal policy responses to changes in the real price of oil. Second, I am interested in studying the robustness of these rules across different specifications in the model space, and in comparing their performance with that of the original Taylor (1993) rule. Therefore, these rules will be employed in the policy analysis in the next section.

The rules were obtained using a grid search of the parameters g_π, g_y, g_i and g_s in (15)

[14] The factor that is commonly used in the model averaging literature (see, for instance, BDW, 2007) to define the set of models with "high" posterior probability is 1/20. The reason why I set the threshold to 1/50 instead of 1/20 is that, in this context, a large number of models, with posterior probability significantly different from zero as a group, do not get captured by the 1/20 threshold. Since I will use the subset of "high" posterior models for the policy evaluation exercise in the next section, the lower threshold of 1/50 allows me to have a group of models that provide a better representation of the model space M.

that minimize the expected loss:

$$\widehat{R}_m = var(\pi_\infty \mid d,p,m) + \lambda_y var(y_\infty \mid d,p,m) + \lambda_i var(\Delta i_\infty \mid d,p,m) \tag{22}$$

for each of the three models described in Table 2.

The rules obtained using this procedure are reported in Table 4. The optimal rule for the BG model implies a stronger response of the nominal interest rate to changes in the real price of oil. As an example, if the real price of oil increases by 10%, the optimal response is to decrease the nominal interest rate by 0.13% if the policymaker believes that H or S are the correct models of the economy, and by 0.52% if he believes that the correct model is BG instead.

Table 4 - Optimal simple rules

	S	BG	H
g_π	1.425	1.625	1.415
g_y	0.377	0.38	0.332
g_i	0.777	0.676	0.788
g_s	-0.013	-0.052	-0.013

Note: Optimal rules in the form described by (15). These rules were obtained by grid search of the coefficients in (15) that minimize (22) for the specification with the highest posterior probability in each class of models.

In order to understand in more detail what these rule imply, and whether the policy recommendations that they suggest are reasonable, I compared them with the actual Federal Funds rate in the period 1960:I - 2008:II. This historical analysis is reported in Figure 2. This figure shows that neither policy rule recommends values for the Federal Funds rate that are constantly and significantly different from the actual pattern of the Federal Funds rate during the period under consideration. In particular, while the BG rule seems to better match the actual Federal Funds rate from the mid 1960s to the mid 1980s, the S and the H rules seem to be closer to the actual pattern of the Federal Funds rate in the time period starting from the second half of the 1980s.

Finally, I study the consequences of the policy responses implied by each of the rules described in Table 4 when the economy is affected by a 10% unexpected increase in the real price of oil. In particular, I investigate the response of output, Core CPI and the Federal Funds rate in each of the models described in Table 2, when the policymaker implements the optimal simple rule for that specific model. For each model, the impact of the policy response to the change in the price of oil is compared with the pattern of the variables of interest when no action is undertaken by the policymaker, i.e. when the coefficients on the policy rule in (15) are all set equal to zero. This exercise provides further evidence about the fact that the ability of monetary policy to contrast an oil price change, and thus the optimal policy response to such change, is model dependent. The results of this exercise are reported in Figure 3.

Some further analysis of the policy responses implied by each of the rules described in Table 4 is provided in Appendix 3.

Figure 2 - Optimal simple rules and true Federal Funds rate: $1960-2008$

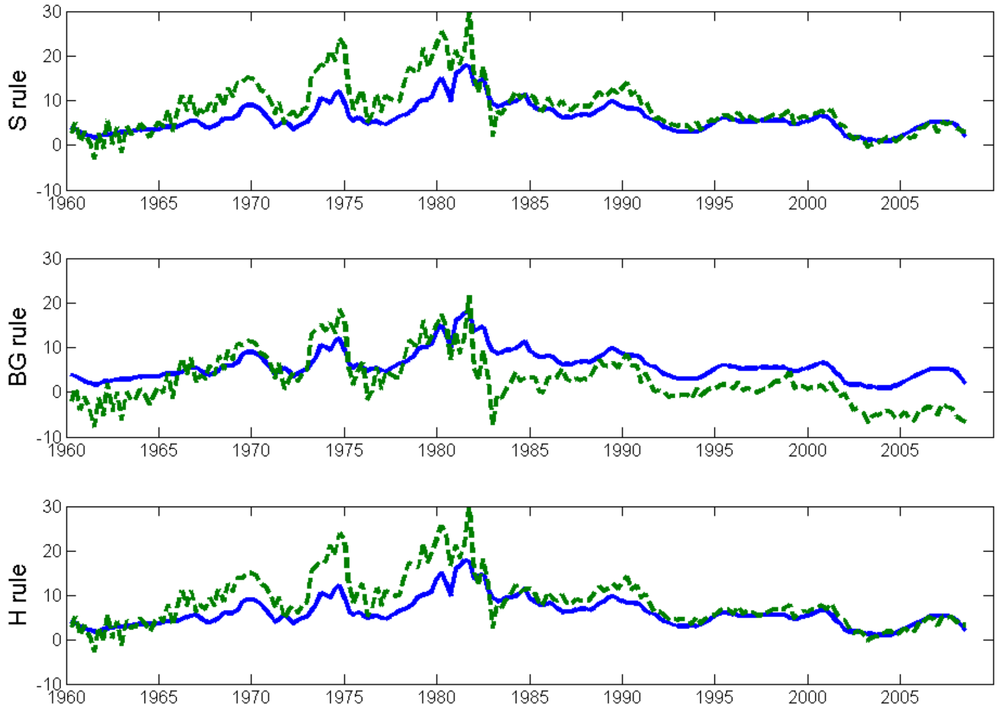

Note: Each panel shows the actual Federal Funds rate (continuous line) and the Federal Funds rate recommended by one of the simple policy rules described in Table 4 (dashed line). The simple policy rules are the S rule in the top panel, the BG rule in the middle panel and the H rule in the bottom panel. The period is 1960:I - 2008:II.

The exercise reported in Figure 3 also provides some additional insight on the debate between Bernanke et al. (1997, 2004) and Hamilton and Herrera (2004) over the ability of monetary policy to avoid the downfalls in output that followed most of the oil price shocks in the postwar period. Bernanke et al. (1997, 2004) argue that the economic downturns would have been milder if the policymaker had adopted a less contractionary policy after an oil price shock. Hamilton and Herrera (2004), on the other hand, argue that output would have decreased anyway, even if the policymaker had kept the Federal Funds rate from increasing.

Figure 3 reports an impulse response exercise that is very similar to those studied by Bernanke et al. (1997, 2004) and Hamilton and Herrera (2004). The different panels in this picture are consistent with the results in both Bernanke et al. (1997, 2004) and Hamilton and Herrera (2004). Thus, this exercise provides evidence that both positions can be correct, depending on which model of the economy is assumed to be the true one.

Figure 3 - Impulse response: optimal simple rules and no action

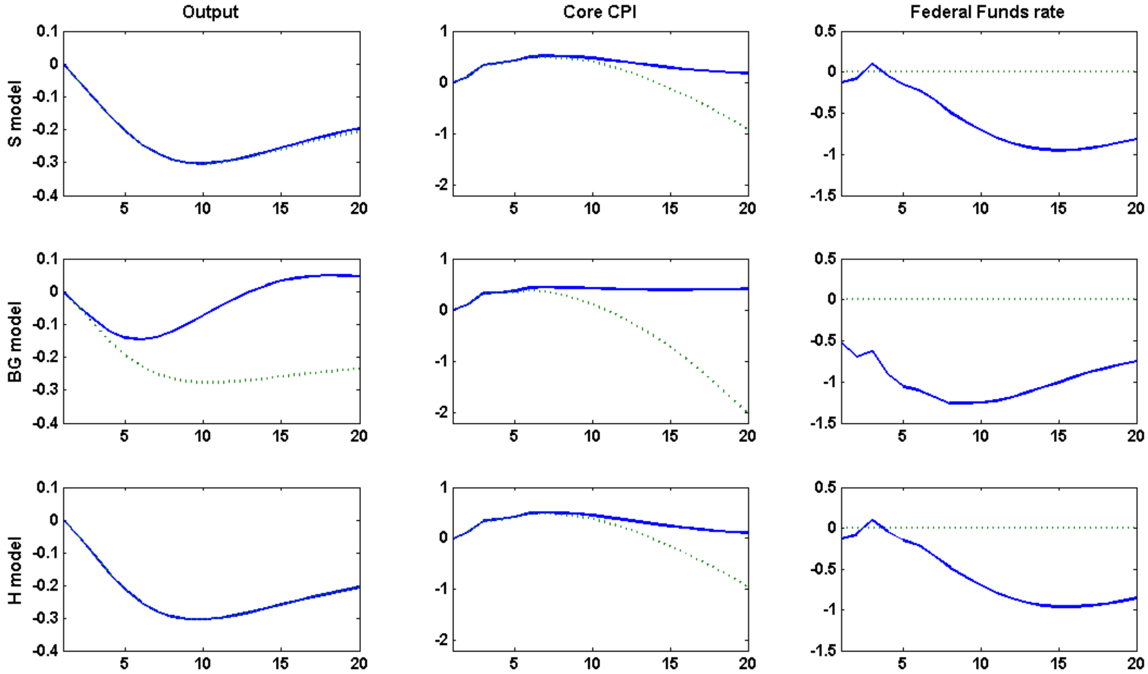

Notes: Response of the output gap, core CPI and the Federal funds rate to a 10% unexpected increase in the real price of oil. The first columns reports the output gap, the second column reports core CPI and the last column reports the Federal Funds rate. Each row represents a different model and policy rule. In each panel, the response of the variable of interest under the selected policy rule (continuous line) is compared to the response when no action is undertaken by the policymaker, i.e. when the coefficients in (15) are all set equal to zero (dashed line).

In more detail, if the true model of the economy is the BG model, then a policymaker that follows the optimal rule reported in Table 4 for this model can successfully avoid the downfall in output caused by the oil price shock. This result for the BG model supports the position of Bernanke et al. (1997, 2004). On the other hand, if the the true model of the economy is the S model or the H model, then the policymaker will not be able to avoid the decrease in output caused by the oil price shock.[15] It follows that in the case of the H model, the response of the policymaker is uniquely directed to stabilize inflation, and in the case of the S model, the response is mainly directed to stabilize inflation, since the response of output is small. For this reason, preventing the Federal Funds rate from increasing would be of small or no benefit, which is the opinion expressed in Hamilton and Herrera (2004).

6 Policy Evaluation

In a context in which the policymaker does not know whether the true model of the economy belongs to the M^S, M^{BG} or M^H class, what are the consequences of adopting a specific policy rule? What rules are more robust across the different specifications? All these questions will be investigated in this section.

The analysis in this section requires the computation of expected losses for the different classes of models in the model space. Expected losses are calculated in the same way as in BDW (2007). Let $\widehat{R}_m$ be the expected loss that occurs conditional on a model and the estimated model parameters, as defined in (22). In addition, let $\widehat{L}_m$ be the BIC adjusted likelihood for model m. Then, for a given set of models, the expected loss when model uncertainty is incorporated in the evaluation is:

$$\widehat{R}_C = \sum_{m \in C} \widehat{R}_m \mu \left(m \mid d \right)$$

Under the assumption of uniform priors, $\mu \left(m \mid d \right)$ is proportional to $\widehat{L}_m$ so we can write:

$$\widehat{R}_C = \frac{\sum_{m \in C} \widehat{R}_m \widehat{L}_m}{\sum_{m \in C} \widehat{L}_m}$$

where, as before, C represents the class of models under analysis. Using this same approach, we can compute the expected loss for the entire model space M.

In the outcome dispersion, action dispersion and minimax analysis that I report next, I restrict the model space to include only the models with "high" posterior probability as defined in Table 3. The reason why I focus on a smaller model space than the one used in the first part of the paper is that some of the specifications in the original model space are unstable, that is they exhibit infinite variance under most of the policy rules under analysis. These specifications have posterior probabilities that are very close to zero, while none of the "high" posterior models exhibits this type of behavior. In order to avoid the policy evaluation

[15] In the H model, the policy is not able to affect the output gap; while in the S model the effect is very small, as shown in Figure 3.

exercise to be driven by models that have insignificant posterior probabilities, I decided to focus on a group of models that exhibits a more stable behavior, and that is able to provide a good representation of the original model space in terms of posterior probability.[16] The model space M and the classes of models M^S, M^{BG} and M^Hwere redefined as a consequence. The new model space includes 294 specifications, while M^S, M^{BG} and M^H are composed of 108, 81 and 105 models respectively. A detailed description of the "high" posterior models used in the policy evaluation exercise, and the definition of the new model space and classes of models are provided in Appendix 1.

6.1 Outcome dispersion

Outcome dispersion measures the variation in loss that occurs when one considers the effects of the same policy rule in different models. I will start by considering outcome dispersion for the standard Taylor (1993) rule, defined as:

$$i_t = 1.5\pi_t + 0.5y_t \tag{23}$$

and for the optimal simple rules obtained in the previous section and described in Table 4.

Table 5 reports the properties of the distribution of losses for each class of models under each of the four policy rules. Figure 4 provides a visual representation of the same results reported in Table 5. This figure clearly shows that in this context the Taylor rule is outperformed, in terms of outcome dispersion, by the alternative simple rules described in Table 4. In particular, while the Taylor rule implies higher and more disperse expected losses in each class of models, it is evident that its performance is significantly worse than the other rules when the true model of the economy is the BG model.

Table 5 - Distribution of model losses under each of the policy rules

(A) Taylor rule	*S*	*BG*	*H*
(1) Mean	17.65	59.42	17.95
(2) Standard deviation	2.76	12.56	2.78
(3) Minimum	13.37	39.46	13.70
(4) Q1	15.07	51.23	15.38
(5) Median	17.04	56.45	17.39
(6) Q3	19.25	66.72	19.50
(7) Maximum	24.42	100.09	24.81
(8) Posterior weighted average	18.11	62.88	18.43
(9) N. of models	108	81	105

[16] See row 7 in Table 3.

(B) S rule	*S*	*BG*	*H*
(1) Mean	12.49	16.40	12.49
(2) Standard deviation	2.92	1.12	2.69
(3) Minimum	10.78	15.18	11.02
(4) Q1	11.44	15.77	11.62
(5) Median	11.60	16.06	11.73
(6) Q3	11.75	16.58	11.90
(7) Maximum	25.82	20.61	26.10
(8) Posterior weighted average	11.91	16.25	12.05
(9) N. of models	108	81	105

(C) BG rule	*S*	*BG*	*H*
(1) Mean	13.42	15.13	13.52
(2) Standard deviation	2.56	1.04	2.35
(3) Minimum	12.04	13.98	12.26
(4) Q1	12.37	14.48	12.58
(5) Median	12.65	14.76	12.84
(6) Q3	13.05	15.33	13.24
(7) Maximum	25.24	18.58	25.50
(8) Posterior weighted average	12.87	14.79	13.04
(9) N. of models	108	81	105

(D) H rule	*S*	*BG*	*H*
(1) Mean	12.45	16.79	12.5
(2) Standard deviation	2.97	1.08	2.73
(3) Minimum	10.82	15.48	11.04
(4) Q1	11.46	16.11	11.63
(5) Median	11.62	16.49	11.73
(6) Q3	11.75	17	11.90
(7) Maximum	26.14	20.68	26.44
(8) Posterior weighted average	11.92	16.53	12.05
(9) N. of models	108	81	105

Note: Distribution of model specific losses for each class of models under the Taylor rule (Panel (A)), the S rule (Panel (B)), the BG rule (Panel (C)) and the H rule (Panel (D)). The composition of each class of models is described in Appendix 1.

Figure 4 - Outcome dispersion for each of the policy rules

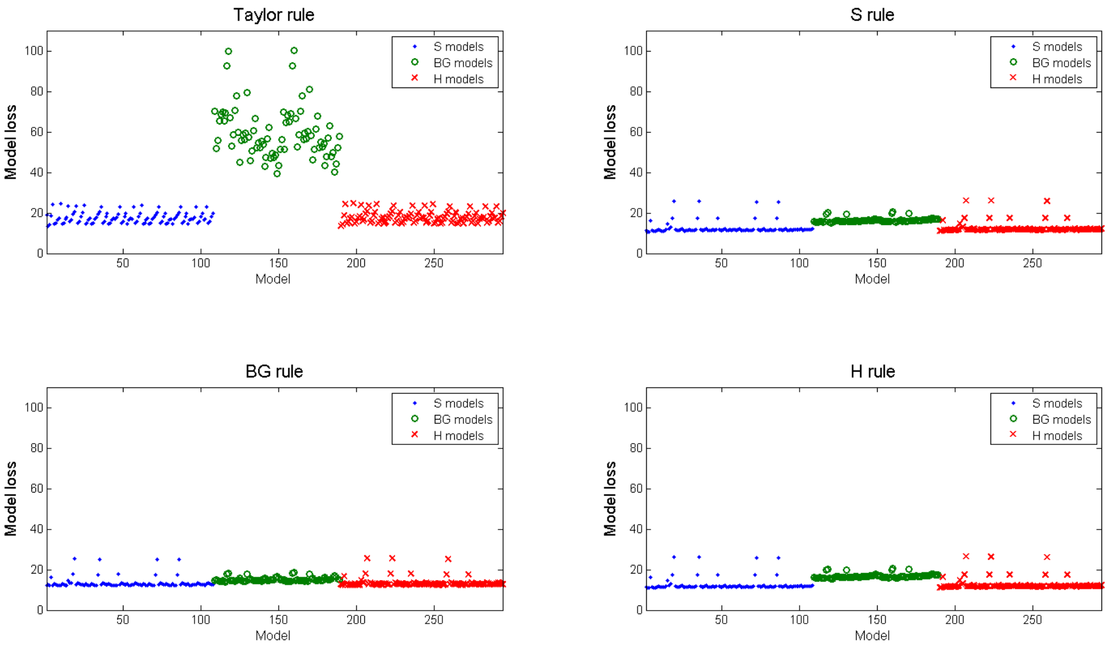

Notes: 1. Each panel reports model specific expected losses under one of the four policies. The Taylor rule is defined in (23), while the S rule, the BG rule and the H rule are described in Table 4.

2. The summary statistics for the distribution of losses in each class of models are reported in Table 5. The summary statistics for the distribution of losses across the model space are reported in Table 6.

3. Model numbers are explained in Appendix 1.

Table 6 - Distribution of model losses in the model space

	Taylor rule	S rule	BG rule	H rule
(1) Mean	29.27	13.55	13.93	13.66
(2) Standard deviation	19.89	3.02	2.28	3.15
(3) Minimum	13.37	10.78	12.04	10.82
(4) Q1	16.05	11.60	12.59	11.61
(5) Median	18.41	11.82	13.05	11.81
(6) Q3	44.82	15.89	14.59	16.23
(7) Maximum	100.09	26.10	25.5	26.44
(8) Posterior weighted average	37.11	13.78	13.73	13.90
(9) N. of models	294	294	294	294

Note: Distribution of model specific losses under the Taylor rule, the S rule, the BG rule and the H rule.

Table 6 reports the distribution of losses across the whole model space for each of the policies under analysis. This table summarizes the results in Table 5, and provides a more clear way of comparing the performance of the different policy rules. In particular, row (8) shows the posterior weighted average loss for each of the policy rules under consideration, which is the measure that is naturally used for policy evaluation in the Bayesian approach.

Figure 5 illustrates the performance of the alternative simple rules relative to the Taylor rule. For each specification, this figure reports the ratio between the loss incurred under each of the simple policy rules and the loss incurred under Taylor rule. As in Figure 4, this figure confirms that each class of models performs better, in average, under the alternative simple rules and that the largest improvement in terms of expected losses is attained by the BG class of models. However, this figure also shows that for some of the S and H specifications, the Taylor rule implies lower losses compared to the alternative simple rules. These specifications are those for which the ratio between the loss under the specific rule and the loss under the Taylor rule is greater than one.

Figure 5. Model losses for each policy relative to the Taylor rule

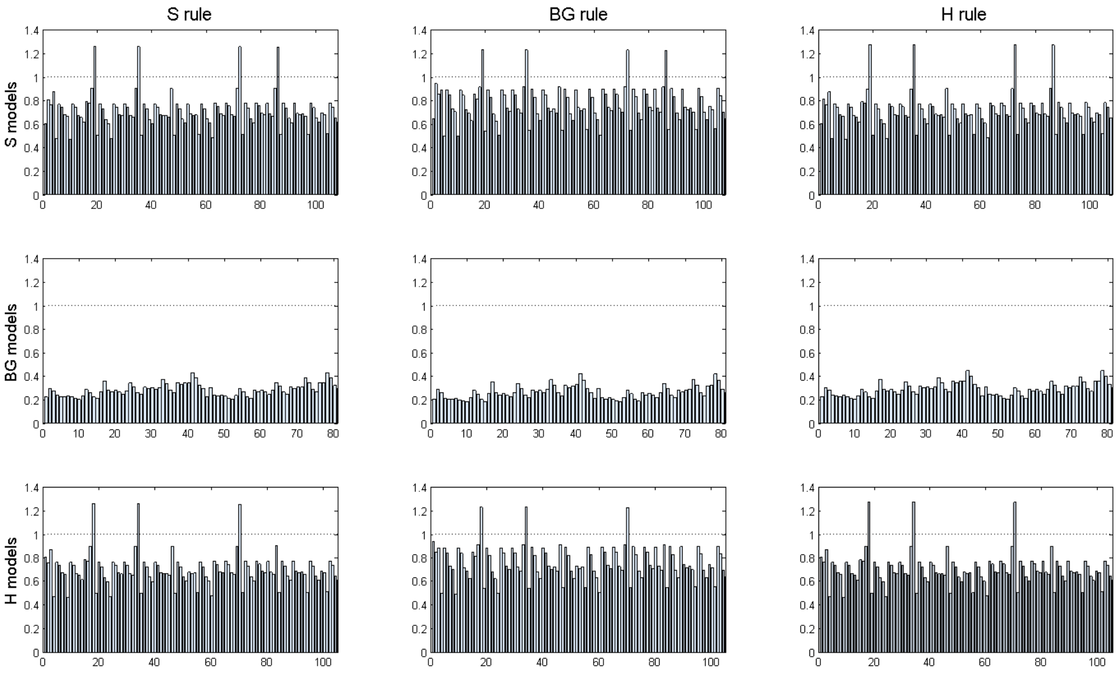

Notes: 1. Each panel reports the performance of a class of models under one of the alternative simple rules relative to the Taylor rule. The specifications for which the Taylor rule implies lower losses are those for which the ratio between the loss under the specific rule and the loss under the Taylor rule is greater than one.

2. Model numbers are explained in Appendix 1.

What is the reason why the Taylor rule performs worse than the alternative simple rules in this context? Figure 6 reports the decomposition of model specific losses into their three components: output gap variance, inflation variance and variance of interest rate differences. This figure shows that the Taylor rule is less successful in controlling the variance of inflation, especially in the *BG* specifications, and this is the reason why it performs worse than the other simple rules in this context.

Figure 6 - Loss decomposition

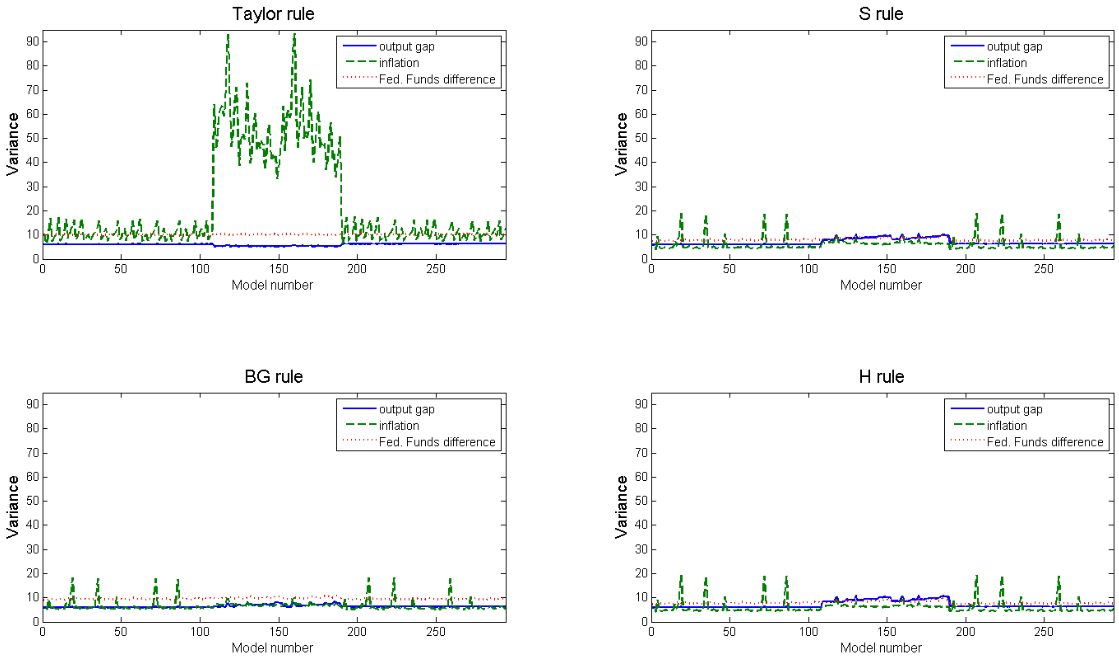

Note: Loss decomposition for the 294 models with "high" posterior probability under the Taylor rule and each of the alternative simple rules. Model numbers are explained in Appendix 1.

6.1.1 The role of inertia in the simple policy rules

Tables $5-6$ and Figures $4-6$ provide definite evidence that in this environment the original Taylor rule performs significantly worse than the simple rules reported in Table 4. Given this result, the policymaker might be interested in investigating whether this happens because the original Taylor rule does not incorporate a term for interest rate smoothing, or because it does not respond to changes in the real price of oil.

To answer this question, I compare the results obtained in the first part of this section with the performance of a Taylor-type rule that includes persistence in the policy instrument. The literature on monetary policy has devoted large attention to the tendency of central banks to adjust interest rates gradually in response to changes in economic conditions. In particular, a number of contributions have focused on the study of inertial Taylor rules. Using the notation adopted in this paper, the typical specification of this type of rule is:

$$i_t = g_i i_{t-1} + (1 - g_i)\widetilde{i}_t \tag{24}$$

where $\widetilde{i}_t$ is the operating target for the policy instrument, and g_i represents the degree of inertia in the central bank's response. The interest rate target is defined as:

$$\widetilde{i}_t = \widetilde{g}_\pi \pi_t + \widetilde{g}_y y_t \tag{25}$$

with $\widetilde{g}_\pi$ and $\widetilde{g}_y$ representing the long run effects of inflation and output gap on the nominal interest rate.

In the next exercise, I will look at the distribution of model losses induced by the policy rule defined in (24) and (25), with $g_i = 0.65$, $\widetilde{g}_\pi = 1.5$ and $\widetilde{g}_y = 0.85$. These values are in the range estimated by the empirical literature on inertial Taylor rules (see for instance Sack, 1998; Orphanides, 2001; Dueker and Rasche, 2004). The policy rule following from these values, which I will call "inertial Taylor rule" (ITR), is $i_t = 0.525\pi_t + 0.2975y_t + 0.65i_{t-1}$.

Table 7 reports the distribution of model losses across the model space for the original Taylor rule (OTR), the ITR, and the three simple rules described in Table 4. This Table shows that the performance of the ITR is similar to that of the OTR, in terms of both outcome dispersion and posterior weighted average loss. It follows that, in this environment, the introduction of inertia in the Taylor rule is not able, per se, to improve the distribution of losses across the specifications in the model space.

Table 7 - Distribution of model losses in the model space

	OTR	ITR	S rule	BG rule	H rule
(1) Mean	29.27	28.7	13.55	13.93	13.66
(2) Standard deviation	19.89	20.87	3.02	2.28	3.15
(3) Minimum	13.37	12.5	10.78	12.04	10.82
(4) Q1	16.05	15.01	11.60	12.59	11.61
(5) Median	18.41	17.15	11.82	13.05	11.81
(6) Q3	44.82	45	15.89	14.59	16.23
(7) Maximum	100.09	103.59	26.10	25.5	26.44
(8) Posterior weighted average	37.11	36.9	13.78	13.73	13.90
(9) N. of models	294	294	294	294	294

Note: Distribution of model losses under different policy rules. The OTR is defined by (23), the ITR is defined by (24) and (25), with g_i= 0.65, $\widetilde{g}_\pi$= 1.5 and $\widetilde{g}_y$= 0.85, the S, BG and H rules are defined by (15), with coefficient values as reported in Table 4.

6.2 Action dispersion

In this section, I consider dispersion over the optimal policy rule solving the policymaker's minimization process. For each specification in the model space, the optimal rule was obtained by grid search of the parameters g_π, g_y, g_i and g_s in (15) that minimize the model specific loss function (22). The action dispersion analysis proposed in this section will focus on the long run effect of the real price of oil on the nominal interest rate, defined as: $g_s/(1-g_i)$.

Table 8 reports some summary statistics of the distribution of $g_s/(1-g_i)$ for each class of models and across the whole model space. Rows (1) and (8) show that, in average, the optimal response to oil price changes is stronger in the BG class of models than in the other two classes. However, the table also shows that the distribution of the optimal policy coefficient $g_s/(1-g_i)$ is significantly more disperse in the H and S classes of models than in the BG class. However, this dispersion seems to be driven by a small subset of models for which the optimal policy response to a change in oil prices is much stronger than for the other specifications in the same class. Finally, the table illustrates that, across the model space, the long run effect of the real price of oil on the nominal interest rate is quite sensitive to the model specification used to compute the optimal policy rule.

Table 8 - Action dispersion results for selected coefficients

	All models		S models		BG models		H models	
	$\frac{g_s}{(1-g_i)}$	g_i	$\frac{g_s}{(1-g_i)}$	g_i	$\frac{g_s}{(1-g_i)}$	g_i	$\frac{g_s}{(1-g_i)}$	g_i
(1) Mean	−0.07	0.73	−0.05	0.77	−0.12	0.63	−0.06	0.78
(2) Standard deviation	0.18	0.09	0.21	0.07	0.03	0.03	0.21	0.06
(3) Minimum	−0.8	0.54	−0.8	0.69	−0.19	0.54	−0.8	0.71
(4) Q1	−0.11	0.67	−0.07	0.72	−0.14	0.6	−0.08	0.74
(5) Median	−0.05	0.73	−0.04	0.77	−0.12	0.63	−0.05	0.78
(6) Q3	0	0.79	0	0.79	−0.1	0.65	0	0.8
(7) Maximum	0.71	0.95	0.57	0.95	−0.05	0.68	0.71	0.95
(8) Posterior weighted average	−0.09	0.72	−0.05	0.76	−0.13	0.64	−0.06	0.78
(9) N. of models	294		108		81		105	

6.3 Minimax and minimax regret

The next two tables report the policy recommendations of the minimax and minimax regret criteria, defined in (6) and (7) in section 1. This exercise is performed using the original Taylor rule and the three simple rules described in Table 4.

Table 9 reports the result of the minimax analysis for each of the model classes and for the whole model space. In each case, the policy recommended by the minimax criterion is the

BG rule. This implies that in each class of models and across the model space, the *BG* rule is the one that minimizes the maximum attainable loss.

Table 9 - Minimax analysis

	(1) All models	(2) S models	(3) BG models	(4) H models
N. of models	294	108	81	105
Maximum				
Taylor rule	100.09	24.42	100.09	24.81
S rule	26.10	25.82	20.61	26.10
BG rule	25.50	25.24	18.58	25.50
H rule	26.44	26.14	20.68	26.44
Minimax	*BG* rule	*BG* rule	*BG* rule	*BG* rule

Table 10 reports the result of the minimax regret criterion for each of the model classes and for the whole model space. Regret is the difference between the loss suffered by a policy relative to the loss under the optimal policy for that specific model, as defined in (8). This table shows that the policy minimizing the maximum regret, in each class of models and in the model space, is the *BG* rule.

Table 10 - Minimax regret analysis

	(1) All models	(2) S models	(3) BG models	(4) H models
N. of models	294	108	81	105
Regret				
Taylor rule	81.51	12.93	81.51	13.14
S rule	5.36	5.30	2.03	5.36
BG rule	4.76	4.73	0	4.76
H rule	5.69	5.63	2.34	5.69
Minimax Regret	*BG* rule	*BG* rule	*BG* rule	*BG* rule

Table 9, column (1) and Table 10, column (1) show that the policy suggested by the minimax and minimax regret analysis are the same, and they correspond to the *BG* rule. From Table 6, row (8) we also know that the policy rule giving the lowest posterior weighted average loss across the 294 models with "high" posterior probability is the *BG* rule. Therefore, in this exercise the policy rule recommended by the Bayesian model averaging approach and the policy rule recommended by the non-Bayesian approaches of the minimax and minimax regret criteria are equivalent.

7 Concluding remarks

In this paper, I studied the problem of a policymaker that is uncertain about the mechanisms through which oil prices affect economic activity. I analyzed three classes of models that have been proposed to explain the effects of oil prices, and I presented a policy evaluation exercise that takes advantage of a range of techniques developed in the model uncertainty literature. I showed that, in this context, the performance of the Taylor rule in terms of loss distribution is significantly worse than the performance of alternative simple rules that introduce persistence in the Federal Funds rate, and that respond to changes in the real price of oil.

The analysis in this paper can be extended in a few different directions. First, the paper shows that model losses are higher, in average, under the Taylor rule because this rule is less successful in controlling the inflation variance component of the loss function. Therefore, I think that it would be appealing to examine how the outcome dispersion results would change if the weights on output and inflation variance in the loss function were allowed to change. In other words, how would the loss dispersion analysis change if the policymaker cared more (or less) about output gap variance than inflation variance?

Second, the policy analysis in this paper could be extended to account for the lack of consensus on the way oil prices should be measured. Indeed, while a part of the literature focuses on real oil prices (Blanchard and Gali, 2007; Herrera and Pesavento, 2007), other contributions refer to different measures of change in the nominal price of oil (Bernanke et al., 1997; Hamilton, 2003). This issue could be incorporated in the framework proposed in this paper by simply considering the uncertainty on the way oil prices are defined as an additional form of uncertainty characterizing the model space.

A third direction would be to include models that focus on allocative disturbances as the channel through which oil prices affect economic activity (see for instance Bernanke, 1983 or Hamilton, 1988). As explained by Hamilton (2004), if this is actually the mechanism through which oil prices affect the economy, then there is no reason to expect a linear relation between oil prices and GDP. An oil price increase will decrease demand for some goods and possibly increase demand for others, and it will create an incentive for households to postpone their investment activity. However, an oil price decrease will have the same effect on the economy, so that both an oil price increase and an oil price decrease could be contractionary in the short run. For this reason, it might be interesting to think about possible ways of including this additional channel of transmission of the effects of oil prices in the policy evaluation exercise.

Finally, a last extension could be in the direction of investigating the role of expectations in this environment. This paper assumes backward expectations; the assumption could be relaxed by using survey data on expected inflation. In addition, it might be interesting to introduce uncertainty on the way expectations are formed, in a way similar to BDW (2007).

Appendix 1

Data description and model labeling

Data description

The variables that I use are the following:

- y_t is output gap, computed as the difference between real GDP and the CBO estimate of potential GDP for any given period, both expressed in logs.
- π_t is the annualized difference in log core CPI, where core CPI is the CPI for all urban consumers: all items less energy products.
- s_t is the real price of oil, defined (as in Blanchard and Gali, 2007) as the difference between the nominal price of oil and core CPI, both expressed in logs. The nominal price of oil is the West Texas Intermediate spot oil price, while core CPI is the same used to compute π_t.
- i_t is the quarterly average Federal Funds rate.

The data is quarterly and includes observations from 1960:I to 2008:II, with data from 1958:IV to 1959:IV used to provide lags. All the data was obtained from the Federal Reserve Bank of St. Louis web site.

Model labeling

The model space includes $3,840$ models, $1,280$ for each class of models. The numbering of the models is organized as follows:

- models from 1 to $1,280$ are the S class of models;
- models from $1,281$ to $2,560$ are the BG class of models;
- models from $2,561$ to $3,840$ are the H class of models.

In each class of models, the specifications differ on the number of lags included in the output and in the inflation equation. In particular, the variables whose lags are allowed to change are described in Table 10:

Table 11 - Description of model specifications

	Output equation		Inflation equation		
	y	s	y	π	s
Lags	$1-4$	$1-4$	$1-4$	$1-4$	$0-4$

The order in which the lags change in the specifications belonging to one class of models is the following:

1. lags of s in the inflation equation;

2. lags of π in the inflation equation;

3. lags of y in the inflation equation;

4. lags of s in the output equation;

5. lags of y in the output equation.

In the policy evaluation exercise, I only consider the models with "high" posterior probability, as defined in Table 3. These models are a subset of the model space, composed of:

- 108 models for the S class;
- 81 models for the BG class;
- 105 models for the H class.

Specifically, the lag composition of each of these models is described in the next tables:

Table 12 - The "high" posterior probability models in the Solow class

		Lags output equation		Lags inflation equation					Lags output equation		Lags inflation equation					Lags output equation		Lags inflation equation		
Number	Original model number	y	s	y	π	s	Number	Original model number	y	s	y	π	s	Number	Original model number	y	s	y	π	s
1	399	2	1	4	4	3	37	797	3	2	4	4	1	73	1036	4	1	4	4	0
2	657	3	1	1	4	1	38	798	3	2	4	4	2	74	1037	4	1	4	4	1
3	659	3	1	1	4	3	39	799	3	2	4	4	3	75	1038	4	1	4	4	2
4	674	3	1	2	3	3	40	800	3	2	4	4	4	76	1039	4	1	4	4	3
5	676	3	1	2	4	0	41	837	3	3	2	4	1	77	1040	4	1	4	4	4
6	677	3	1	2	4	1	42	838	3	3	2	4	2	78	1077	4	2	2	4	1
7	678	3	1	2	4	2	43	839	3	3	2	4	3	79	1078	4	2	2	4	2
8	679	3	1	2	4	3	44	840	3	3	2	4	4	80	1079	4	2	2	4	3
9	680	3	1	2	4	4	45	859	3	3	3	4	3	81	1080	4	2	2	4	4
10	696	3	1	3	4	0	46	860	3	3	3	4	4	82	1097	4	2	3	4	1
11	697	3	1	3	4	1	47	874	3	3	4	3	3	83	1099	4	2	3	4	3
12	698	3	1	3	4	2	48	876	3	3	4	4	0	84	1100	4	2	3	4	4
13	699	3	1	3	4	3	49	877	3	3	4	4	1	85	1114	4	2	4	3	3
14	700	3	1	3	4	4	50	878	3	3	4	4	2	86	1115	4	2	4	3	4
15	711	3	1	4	3	0	51	879	3	3	4	4	3	87	1116	4	2	4	4	0
16	712	3	1	4	3	1	52	880	3	3	4	4	4	88	1117	4	2	4	4	1
17	713	3	1	4	3	2	53	917	3	4	2	4	1	89	1118	4	2	4	4	2
18	714	3	1	4	3	3	54	919	3	4	2	4	3	90	1119	4	2	4	4	3
19	715	3	1	4	3	4	55	920	3	4	2	4	4	91	1120	4	2	4	4	4
20	716	3	1	4	4	0	56	939	3	4	3	4	3	92	1157	4	3	2	4	1
21	717	3	1	4	4	1	57	956	3	4	4	4	0	93	1159	4	3	2	4	3
22	718	3	1	4	4	2	58	957	3	4	4	4	1	94	1160	4	3	2	4	4
23	719	3	1	4	4	3	59	958	3	4	4	4	2	95	1179	4	3	3	4	3
24	720	3	1	4	4	4	60	959	3	4	4	4	3	96	1180	4	3	3	4	4
25	756	3	2	2	4	0	61	960	3	4	4	4	4	97	1196	4	3	4	4	0
26	757	3	2	2	4	1	62	996	4	1	2	4	0	98	1197	4	3	4	4	1
27	758	3	2	2	4	2	63	997	4	1	2	4	1	99	1198	4	3	4	4	2
28	759	3	2	2	4	3	64	998	4	1	2	4	2	100	1199	4	3	4	4	3
29	760	3	2	2	4	4	65	999	4	1	2	4	3	101	1200	4	3	4	4	4
30	777	3	2	3	4	1	66	1000	4	1	2	4	4	102	1239	4	4	2	4	3
31	778	3	2	3	4	2	67	1017	4	1	3	4	1	103	1240	4	4	2	4	4
32	779	3	2	3	4	3	68	1018	4	1	3	4	2	104	1276	4	4	4	4	0
33	780	3	2	3	4	4	69	1019	4	1	3	4	3	105	1277	4	4	4	4	1
34	794	3	2	4	3	3	70	1020	4	1	3	4	4	106	1278	4	4	4	4	2
35	795	3	2	4	3	4	71	1034	4	1	4	3	3	107	1279	4	4	4	4	3
36	796	3	2	4	4	0	72	1035	4	1	4	3	4	108	1280	4	4	4	4	4

Table 13 - The "high" posterior probability models in the Blanchard-Gali class

		Lags output equation		Lags inflation equation					Lags output equation		Lags inflation equation					Lags output equation		Lags inflation equation		
Number	Original model number	y	s	y	π	s	Number	Original model number	y	s	y	π	s	Number	Original model number	y	s	y	π	s
109	1956	3	1	2	4	0	136	2119	3	3	2	4	3	163	2318	4	1	4	4	2
110	1957	3	1	2	4	1	137	2120	3	3	2	4	4	164	2319	4	1	4	4	3
111	1958	3	1	2	4	2	138	2139	3	3	3	4	3	165	2320	4	1	4	4	4
112	1959	3	1	2	4	3	139	2140	3	3	3	4	4	166	2359	4	2	2	4	3
113	1960	3	1	2	4	4	140	2156	3	3	4	4	0	167	2360	4	2	2	4	4
114	1976	3	1	3	4	0	141	2157	3	3	4	4	1	168	2379	4	2	3	4	3
115	1979	3	1	3	4	3	142	2158	3	3	4	4	2	169	2380	4	2	3	4	4
116	1980	3	1	3	4	4	143	2159	3	3	4	4	3	170	2394	4	2	4	3	3
117	1994	3	1	4	3	3	144	2160	3	3	4	4	4	171	2396	4	2	4	4	0
118	1995	3	1	4	3	4	145	2199	3	4	2	4	3	172	2397	4	2	4	4	1
119	1996	3	1	4	4	0	146	2200	3	4	2	4	4	173	2398	4	2	4	4	2
120	1997	3	1	4	4	1	147	2219	3	4	3	4	3	174	2399	4	2	4	4	3
121	1998	3	1	4	4	2	148	2236	3	4	4	4	0	175	2400	4	2	4	4	4
122	1999	3	1	4	4	3	149	2237	3	4	4	4	1	176	2439	4	3	2	4	3
123	2000	3	1	4	4	4	150	2238	3	4	4	4	2	177	2440	4	3	2	4	4
124	2036	3	2	2	4	0	151	2239	3	4	4	4	3	178	2459	4	3	3	4	3
125	2037	3	2	2	4	1	152	2240	3	4	4	4	4	179	2476	4	3	4	4	0
126	2039	3	2	2	4	3	153	2276	4	1	2	4	0	180	2477	4	3	4	4	1
127	2040	3	2	2	4	4	154	2277	4	1	2	4	1	181	2478	4	3	4	4	2
128	2059	3	2	3	4	3	155	2279	4	1	2	4	3	182	2479	4	3	4	4	3
129	2060	3	2	3	4	4	156	2280	4	1	2	4	4	183	2480	4	3	4	4	4
130	2074	3	2	4	3	3	157	2299	4	1	3	4	3	184	2519	4	4	2	4	3
131	2076	3	2	4	4	0	158	2300	4	1	3	4	4	185	2556	4	4	4	4	0
132	2077	3	2	4	4	1	159	2314	4	1	4	3	3	186	2557	4	4	4	4	1
133	2078	3	2	4	4	2	160	2315	4	1	4	3	4	187	2558	4	4	4	4	2
134	2079	3	2	4	4	3	161	2316	4	1	4	4	0	188	2559	4	4	4	4	3
135	2080	3	2	4	4	4	162	2317	4	1	4	4	1	189	2560	4	4	4	4	4

Table 14 - The "high" posterior probability models in the Hamilton class

		Lags output equation		Lags inflation equation					Lags output equation		Lags inflation equation					Lags output equation		Lags inflation equation		
Number	Original model number	y	s	y	π	s	Number	Original model number	y	s	y	π	s	Number	Original model number	y	s	y	π	s
190	3217	3	1	1	4	1	225	3357	3	2	4	4	1	260	3596	4	1	4	4	0
191	3219	3	1	1	4	3	226	3358	3	2	4	4	2	261	3597	4	1	4	4	1
192	3234	3	1	2	3	3	227	3359	3	2	4	4	3	262	3598	4	1	4	4	2
193	3236	3	1	2	4	0	228	3360	3	2	4	4	4	263	3599	4	1	4	4	3
194	3237	3	1	2	4	1	229	3397	3	3	2	4	1	264	3600	4	1	4	4	4
195	3238	3	1	2	4	2	230	3398	3	3	2	4	2	265	3637	4	2	2	4	1
196	3239	3	1	2	4	3	231	3399	3	3	2	4	3	266	3638	4	2	2	4	2
197	3240	3	1	2	4	4	232	3400	3	3	2	4	4	267	3639	4	2	2	4	3
198	3256	3	1	3	4	0	233	3419	3	3	3	4	3	268	3640	4	2	2	4	4
199	3257	3	1	3	4	1	234	3420	3	3	3	4	4	269	3657	4	2	3	4	1
200	3258	3	1	3	4	2	235	3434	3	3	4	3	3	270	3659	4	2	3	4	3
201	3259	3	1	3	4	3	236	3436	3	3	4	4	0	271	3660	4	2	3	4	4
202	3260	3	1	3	4	4	237	3437	3	3	4	4	1	272	3674	4	2	4	3	3
203	3271	3	1	4	3	0	238	3438	3	3	4	4	2	273	3676	4	2	4	4	0
204	3272	3	1	4	3	1	239	3439	3	3	4	4	3	274	3677	4	2	4	4	1
205	3273	3	1	4	3	2	240	3440	3	3	4	4	4	275	3678	4	2	4	4	2
206	3274	3	1	4	3	3	241	3479	3	4	2	4	3	276	3679	4	2	4	4	3
207	3275	3	1	4	3	4	242	3480	3	4	2	4	4	277	3680	4	2	4	4	4
208	3276	3	1	4	4	0	243	3499	3	4	3	4	3	278	3717	4	3	2	4	1
209	3277	3	1	4	4	1	244	3516	3	4	4	4	0	279	3719	4	3	2	4	3
210	3278	3	1	4	4	2	245	3517	3	4	4	4	1	280	3720	4	3	2	4	4
211	3279	3	1	4	4	3	246	3518	3	4	4	4	2	281	3739	4	3	3	4	3
212	3280	3	1	4	4	4	247	3519	3	4	4	4	3	282	3740	4	3	3	4	4
213	3316	3	2	2	4	0	248	3520	3	4	4	4	4	283	3756	4	3	4	4	0
214	3317	3	2	2	4	1	249	3556	4	1	2	4	0	284	3757	4	3	4	4	1
215	3318	3	2	2	4	2	250	3557	4	1	2	4	1	285	3758	4	3	4	4	2
216	3319	3	2	2	4	3	251	3558	4	1	2	4	2	286	3759	4	3	4	4	3
217	3320	3	2	2	4	4	252	3559	4	1	2	4	3	287	3760	4	3	4	4	4
218	3337	3	2	3	4	1	253	3560	4	1	2	4	4	288	3799	4	4	2	4	3
219	3338	3	2	3	4	2	254	3577	4	1	3	4	1	289	3800	4	4	2	4	4
220	3339	3	2	3	4	3	255	3578	4	1	3	4	2	290	3836	4	4	4	4	0
221	3340	3	2	3	4	4	256	3579	4	1	3	4	3	291	3837	4	4	4	4	1
222	3354	3	2	4	3	3	257	3580	4	1	3	4	4	292	3838	4	4	4	4	2
223	3355	3	2	4	3	4	258	3594	4	1	4	3	3	293	3839	4	4	4	4	3
224	3356	3	2	4	4	0	259	3595	4	1	4	3	4	294	3840	4	4	4	4	4

Appendix 2

Estimating the process for the real price of oil

Given the econometric model described in (18) and (19), the estimates for the mean and variance of d_t are defined as:

$$\widehat{d}_{t|t-1} \equiv E\left(d_t \mid \Im_{t-1}\right)$$

$$P_{t|t-1} \equiv Var\left(d_t \mid \Im_{t-1}\right)$$

$$\Im_t \equiv \{y_1, \pi_1, ..., y_t, \pi_t\}$$

Given the initial $\widehat{d}_{1|0}$ and $P_{1|0}$, I assume that the policymaker updates his estimates using the Kalman filter algorithm:

$$\widehat{d}_{t+1|t} = \widehat{d}_{t|t-1} + \frac{P_{t|t-1}\Phi_t\left(s_t - \Phi_t'\widehat{d}_{t|t-1}\right)}{\sigma_o^2 + \Phi_t' P_{t|t-1}\Phi_t}$$

$$P_{t+1|t} = P_{t|t-1} - \frac{P_{t|t-1}\Phi_t\Phi_t' P_{t|t-1}}{\sigma_o^2 + \Phi_t' P_{t|t-1}\Phi_t} + V$$

The values of $\widehat{d}_{1|0}$ and $P_{1|0}$ were estimated using data from 1957:I to 1969:IV. The initial beliefs on the coefficients of the process for the real price of oil were: $\widehat{d}_{1|0} = [2.041 \quad 0.975]'$. The standard deviation of the estimate of the initial value of ρ is 0.0223, which implies that using data up to the last quarter in 1969, the policymaker cannot reject the null hypothesis of a unit root process for the real price of oil.

The variance of the oil price shock σ_o^2 and of the variance matrix V were set as: $\sigma_o^2 = (13.2)^2$ and

$$V = \begin{bmatrix} (6.3)^2 & 0 \\ 0 & 0 \end{bmatrix}$$

These values can be interpreted as the beliefs of the policymaker over the variance of the oil price shocks, and over the variance of the drift in the mean of the process for the real price of oil. Although these values are arbitrary, they have been set so that the conditional variance of s_t matches the value estimated in Blanchard and Gali (2007).[17] I experimented with different combinations of σ_o^2 and σ_ε^2, and the resulting estimates of ρ are not very sensitive to changes in these values.

[17] See Table 4 in Blanchard and Gali (2007). Specifically, the values of σ_0^2 and V have been chosen to be consistent with the average conditional variance estimates for the sample periods 1960:I - 1983:IV and 1984:I - 2005:IV.

Appendix 3

Analysis of the policy responses implied by the rules described in Table 4

This Appendix provides a more in depth analysis of the policy responses implied by the optimal simple rules reported in Table 4.

Figure 7 investigates the response of output, Core CPI and the Federal Funds rate to a 10% increase in the real price of oil in each of the models described in Table 2, when the policymaker implements the optimal simple rule for that specific model. For each model, the impact of the policy response to the change in the price of oil is compared with the pattern of the variables of interest under the original Taylor rule described by (23).

Figure 7 - Impulse response: optimal simple rules and Taylor rule

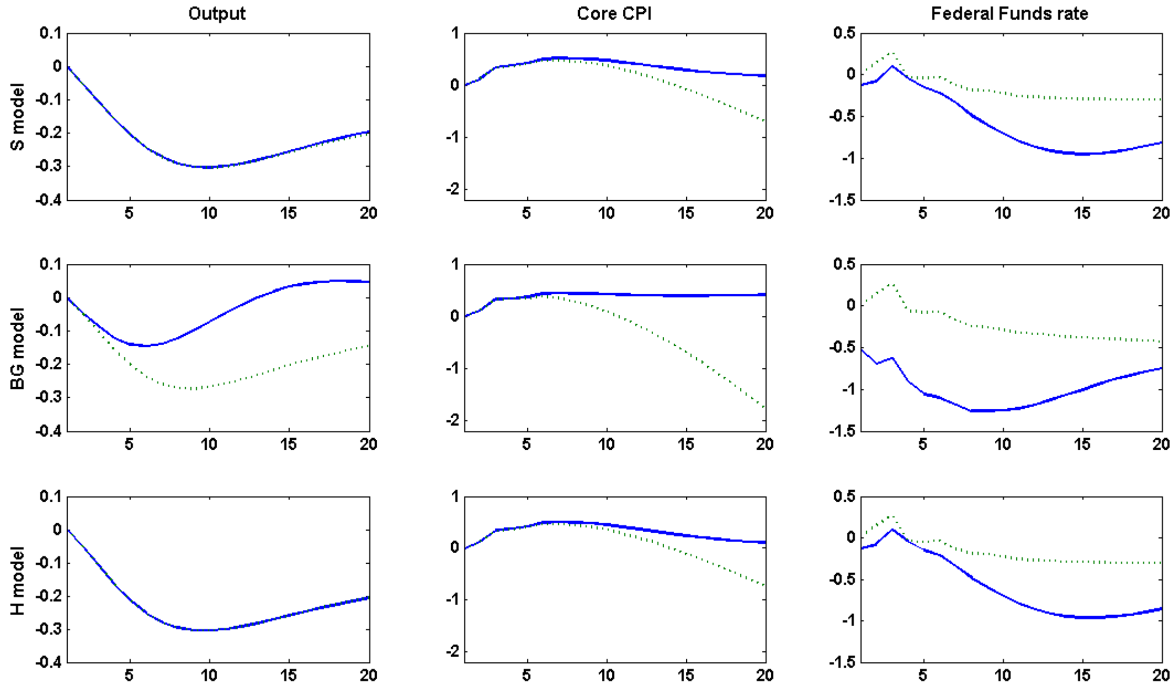

Notes: Response of the output gap, core CPI and the Federal funds rate to a 10% unexpected increase in the real price of oil. The first columns reports the output gap, the second column reports core CPI and the last column reports the Federal Funds rate. Each row represents a different model and policy rule. In each panel, the response of the variable of interest under the optimal simple rule (continuous line) is compared to the response under the original Taylor rule described by (23) (dashed line).

The last column of Figure 7 provides information on the optimal response of the Federal funds rate for each of the models under analysis. In the S model and in the H model, the optimal simple rule implies a less contractionary policy response than the original Taylor rule. On the other hand, in the *BG* model the optimal simple rule implies an initial decrease of the Federal Funds rate, thus recommending an expansionary policy response to a 10% increase in the real price of oil.

The next three figures explore the response of output, Core CPI and the Federal Funds rate to a 10% increase in the real price of oil in each of the models described in Table 2, when the policymaker implements one of the simple policy rules described in Table 4. Figure 8 compares the pattern of the variables of interest under the S rule to their behavior under the original Taylor rule described by (23).

In all the three models, the S rule implies a less contractionary response of the Federal Funds rate, and a smaller decrease in Core CPI than the original Taylor rule. In the *BG* model, and to a small extent also in the S model, the S rule will also be able to reduce the decrease in output caused by the increase in the real price of oil.

Figure 8 - Impulse response: S rule and Taylor rule

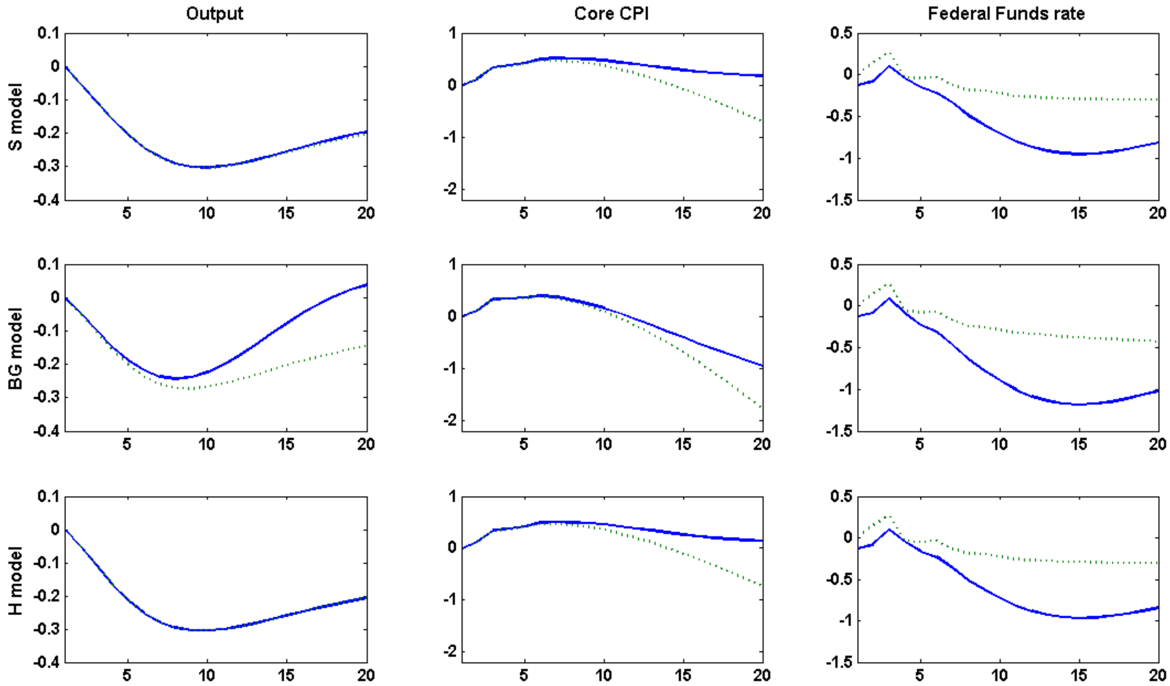

Notes: Response of the output gap, core CPI and the Federal funds rate to a 10% unexpected increase in the real price of oil. The first columns reports the output gap, the second column reports core CPI and the last column reports the Federal Funds rate. Each row represents a different model. In each panel, the response of the variable of interest under the S policy rule (continuous line) is compared to the response under the original Taylor rule described by (23) (dashed line).

Figure 9 compares the pattern of the variables of interest under the BG rule to their behavior under the original Taylor rule described by (23). This Figure shows that the stronger response to a change in the real price of oil recommended by BG rule implies a decrease in the Federal Funds rate in all the models described in Table 2.

Figure 9 - Impulse response: BG rule and Taylor rule

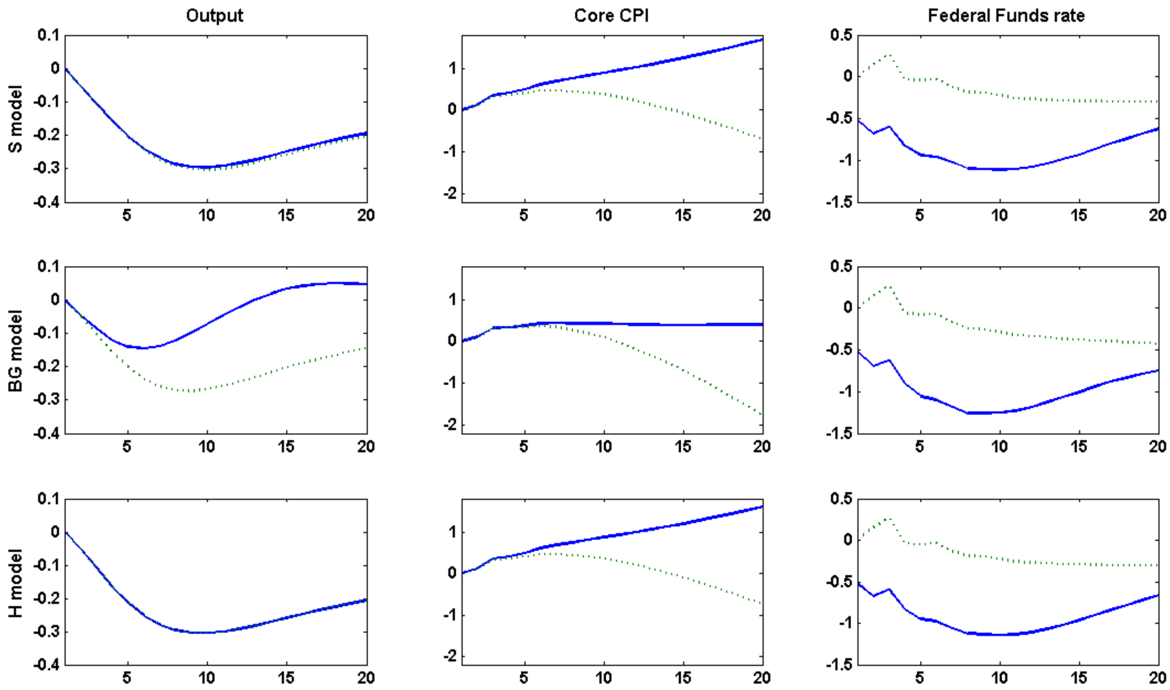

Notes: Response of the output gap, core CPI and the Federal funds rate to a 10% unexpected increase in the real price of oil. The first columns reports the output gap, the second column reports core CPI and the last column reports the Federal Funds rate. Each row represents a different model. In each panel, the response of the variable of interest under the BG policy rule (continuous line) is compared to the response under the original Taylor rule described by (23) (dashed line).

Finally, Figure 10 compares the pattern of the variables of interest under the H rule to their behavior under the original Taylor rule described by (23). Given that the H rule and the S rule are very similar, they will imply an almost identical response of output, Core CPI and the Federal Funds rate to a 10% increase in the real price of oil, as evident from comparing this Figure with Figure 8.

Figure 10 - Impulse response: H rule and Taylor rule

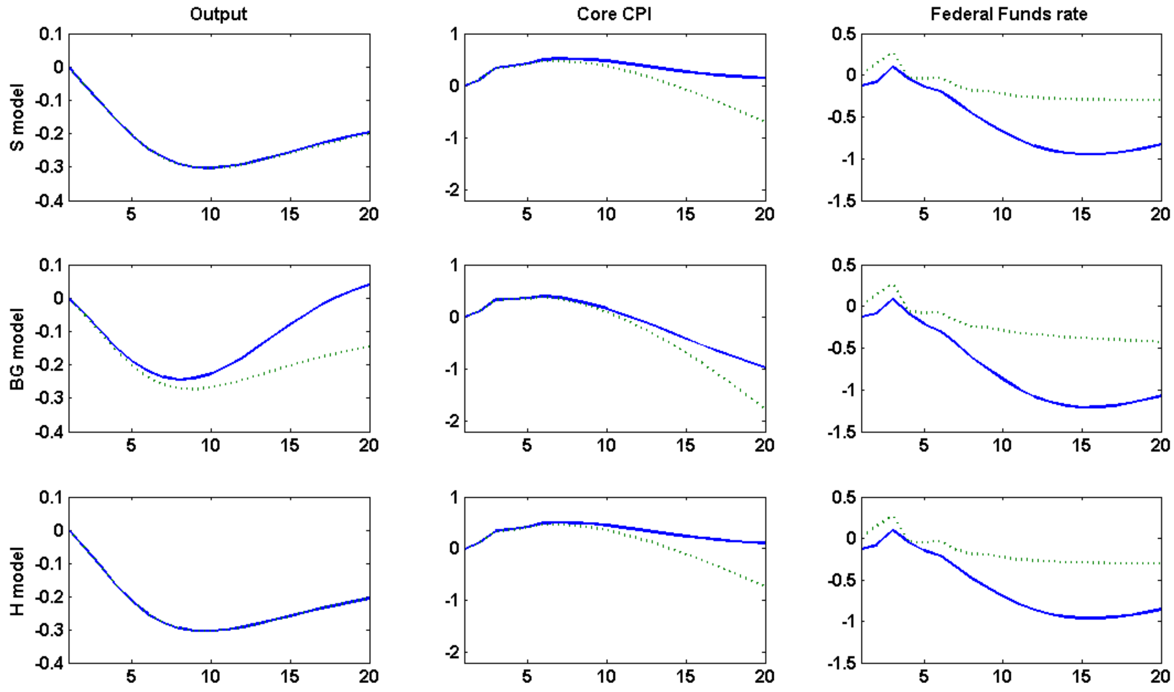

Notes: Response of the output gap, core CPI and the Federal funds rate to a 10% unexpected increase in the real price of oil. The first columns reports the output gap, the second column reports core CPI and the last column reports the Federal Funds rate. Each row represents a different model. In each panel, the response of the variable of interest under the H policy rule (continuous line) is compared to the response under the original Taylor rule described by (23) (dashed line).

References

[1] Bernanke, B.S., 1983. Irreversibility, uncertainty, and cyclical investment. Quarterly Journal of Economics 98, 85–106.

[2] Bernanke, B. S., Gertler, M., Watson, M., 1997. Systematic monetary policy and the effects of oil price shocks. Brookings Papers on Economic Activity, no. 1, pp. 91-142.

[3] Bernanke, B. S., Gertler, M., Watson, M., 2004. Oil shocks and aggregate macroeconomic behavior: the role of monetary policy: reply. Journal of Money, Credit, and Banking, vol. 36, no. 2, April, pp. 287-91.

[4] Blanchard, O. J., Gali, J., 2007. The Macroeconomic Effects of Oil Shocks: Why are the 2000s So Different from the 1970s? National Bureau of Economic Research, Working Papers: 13368.

[5] Brock, W. A., Durlauf, S. N., 2004. Elements of a Theory of Design Limits to Optimal Policy. NBER Working Papers: 10495.

[6] Brock, W. A., Durlauf, S. N., Nason, J. M., Rondina, G., 2007. Simple versus Optimal Rules as Guides to Policy. Journal of Monetary Economics, vol. 54, no. 5, July, pp. 1372-96.

[7] Brock, W. A., Durlauf, S. N, West, K. D., 2003. Policy Evaluation in Uncertain Economic Environments. Brookings Papers on Economic Activity, pp.235-301.

[8] Brock, W. A., Durlauf, S. N, West, K. D., 2007. Model Uncertainty and Policy Evaluation: Some Theory and Empirics. Journal of Econometrics, vol. 136, no. 2, February, pp. 629-64.

[9] Dueker, M.J., Rasche, R.H., 2004. Discrete Policy Changes and Empirical Models of the Federal Funds Rate. Federal Reserve Bank of St. Louis Review, vol. 86, no. 6, November-December, pp. 61-72.

[10] Cogley, T., Colacito, R., Sargent, T. J., 2007. Benefits from U.S. Monetary Policy Experimentation in the Days of Samuelson and Solow and Lucas. Journal of Money, Credit, and Banking, vol. 39, Supplement February, pp. 67-99.

[11] Cogley, T., Sargent, T. J., 2005. The Conquest of U.S. inflation: Learning and Robustness to Model Uncertainty. Review of Economic Dynamics, vol. 8, no. 2, April, pp. 262-302.

[12] Dow, S. C., 2004. Uncertainty and Monetary Policy. Oxford Economic Papers, vol. 56, no. 3, July, pp. 539-61.

[13] Finn, M.G., 2000. Perfect competition and the effects of energy price increases on economic activity. Journal of Money, Credit, and Banking 32, 400–416.

[14] Giannoni, M. P., 2007. Robust Optimal Monetary Policy in a Forward-Looking Model with Parameter and Shock Uncertainty. Journal of Applied Econometrics, vol. 22, no. 1, January-February, pp. 179-213.

[15] Hamilton, J.D., 1988. A neoclassical model of unemployment and the business cycle. Journal of Political Economy 96, 593–617.

[16] Hamilton, J. D., 2003. What is an oil price shock? Journal of Econometrics, vol. 113, pp. 363-398.

[17] Hamilton, J.D., 2005. 'Oil and the Macroeconomy', mimeo, Department of Economics, University of California, San Diego, prepared for Palgrave Dictionary of Economics.

[18] Hamilton, J. D., Herrera A. M., 2004. Oil shocks and aggregate macroeconomic behavior: the role of monetary policy: comment. Journal of Money, Credit, and Banking, vol. 36, no. 2, April, pp. 265-86.

[19] Hansen, L., Sargent, T., 2001a. Robust control and model uncertainty. American Economic Review, vol. 91, no. 2, May, pp. 60-66.

[20] Hansen, L., Sargent, T., 2001b. Acknowledging Misspecification in Macroeconomic Theory. Review of Economic Dynamics, vol. 4, no. 3, July, pp. 519-35.

[21] Hansen, L., Sargent, T., 2008. Robustness. Princeton University Press, Princeton, New Jersey.

[22] Herrera, A. M., Pesavento, E., 2007. Oil Price Shocks, Systematic Monetary Policy and the "Great Moderation". Mimeo, Emory University.

[23] Kim, I., Loungani, P., 1992. The role of energy in real business cycle models. Journal of Monetary Economics 29, 173–189.

[24] Leduc, S., Sill, K., 2004. A quantitative analysis of oil price shocks, systematic monetary policy and economic downturns. Journal of Monetary Economics, vol. 51, pp. 781-808.

[25] Levin, A., Williams, J., 2003. Robust monetary policy with competing reference models. Journal of Monetary Economics 50, 945–975.

[26] Mork, K. A., 1994. Business cycles and the oil market (special issue). Energy Journal 15, pp. 15–37.

[27] Onatski, A., Stock, J., 2002. Robust monetary policy under model uncertainty in a small model of the US economy. Macroeconomic Dynamics 6, pp. 85–110.

[28] Orphanides, A., 2001. Monetary Policy Rules Based on Real-Time Data. American Economic Review, vol. 91, no. 4, September, pp. 964-85.

[29] Pindyck, R.S., 1980. Energy price increases and macroeconomic policy. Energy Journal 1, 1–20.

[30] Rasche, R.H., Tatom, J.A., 1977. Energy resources and potential GNP. Federal Reserve Bank of St. Louis Review 59 (June), 10–24.

[31] Rotemberg, J.J., Woodford, M., 1996. Imperfect competition and the effects of energy price increases. Journal of Money, Credit, and Banking 28, 549–577.

[32] Raftery, Adrian E. 1995. Bayesian Model Selection in Social Research. Sociological Methodology 25: 111–95.

[33] Rudebusch, G., Svensson, L., 1999. Policy rules for inflation targeting. In: Taylor, J. (Ed.), Monetary Policy Rules. University of Chicago Press, Chicago, pp. 203–246.

[34] Sack, B., 1998. Uncertainty, learning, and gradual monetary policy. Board of Governors of the Federal Reserve System, Finance and Economics Discussion Series: 1998-34.

[35] Sargent, T. J., Williams, N., 2005. Impacts of Priors on Convergence and Escapes from Nash Inflation. Review of Economic Dynamics, vol. 8, n. 2, pp. 360–91.

[36] Segal, P., 2007. Why do oil price shocks no longer shock? Oxford Institute for Energy Studies Working Paper.

[37] Solow, R. M., 1980. What to do (macroeconomically) when OPEC comes. In: Fischer, S. (Ed.), Rational Expectations and Economic Policy. University of Chicago Press, Chicago.

[38] Taylor, J., 1993. Discretion versus policy rules in practice. Carnegie-Rochester Conference Series on Public Policy 39, 195–214.

[39] Woodford, M., 2003. Interest and prices. Princeton and Oxford: Princeton University Press, pp. xv, 785.