Policy Evaluation and Uncertainty about the Effects of Oil Prices on Economic Activity

Francesca Rondina

November 2010

Barcelona Economics Working Paper Series

Working Paper nº 522

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

Francesca Rondina*†

Institute for Economic Analysis, CSIC and Barcelona GSE

November 2010

Abstract

This paper addresses the issue of policy evaluation in a context in which policymakers are uncertain about the effects of oil prices on economic performance. I consider models of the economy inspired by Solow (1980), Blanchard and Gali (2007), Kim and Loungani (1992) and Hamilton (1983, 2005), which incorporate different assumptions on the channels through which oil prices have an impact on economic activity. I first study the characteristics of the model space and I analyze the likelihood of the different specifications. I show that the existence of plausible alternative representations of the economy forces the policymaker to face 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 find that, in the environment under analysis, the standard Taylor rule is outperformed under a number of criteria by alternative simple rules in which policymakers introduce persistence in the policy instrument and respond to changes in the real price of oil.

JEL Classification: C52, E52, E58

Keywords: model uncertainty, robust policy, Bayesian model averaging, minimax, oil prices.

*Contacts: Campus UAB, 08193 Bellaterra, Barcelona, Spain; tel: +34 935 806 612, fax: +34 935 801 452, email: francesca.rondina@iae.csic.es

†I am especially grateful to Steven Durlauf for his guidance and constant encouragement. I am indebted to William Brock, Giacomo Rondina and Emmanuele Bobbio for valuable discussions and suggestions. I have also benefited from comments from Kenneth West, Noah Williams, Federico Diez, Nonarit Bisonyabut and seminar participants at the Bank of England, Riksbank, Pomona College, Indiana University, Universidade Nova, BIS, EPFL-UNIL, Dallas Fed, Vassar College and IAE-CSIC. All errors remain my own. Financial support from the Government of Catalonia and the Spanish Ministry for Science and Innovation (Programa Operativo FSE 2007-2013) is gratefully acknowledged.

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 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 a wide range of measures, based on a number of different approaches, 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

[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 the uncertainty on the way the public forms expectations on future economic variables, while in Cogley and Sargent (2005) policymakers are uncertain about the specification of Phillips curve to be adopted for policy decisions.

optimally respond to changes in oil prices.[3] Hence, this is a context in which the application of the 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 a change in oil prices. 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 work 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 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 main frameworks that have been proposed to explain the effects of oil prices on economic variables. For each of these frameworks, I study the consequences of the implementation of alternative simple policy rules, and I investigate the extent to which 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 the specific form of uncertainty under analysis in this paper.

This study 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 policy 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 elements of the model space do not necessarily lie within

[3] An example of this disagreement is the debate between Bernanke, Gertler and Watson (1997, 2004) and 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 minimizes while nature maximize losses 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 et al. (2010) are examples of applications of this approach to the analysis of monetary policy.

some distance from a baseline model specification.

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 description of the model space in sections 4 will provide more 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 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 of the characteristics of the model space.

This work is also related to the large literature studying the impact of oil prices on economic activity. This paper does not intend to take a position in the debate over the different models proposed to explain the effects of a change in oil prices on economic performance. Rather, I show that different frameworks, based on different channels of transmission of oil price shocks into the economy, are plausible alternative approximations of the true data generating process. Finally, this paper is related to the literature investigating 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 on its contribution to the milder reaction of economic variables to oil price shocks since the mid 1980s (Blanchard and Gali, 2007; Herrera and Pesavento, 2009; Clark and Terry, 2010). This work provides some additional insights in this area by explicitly analyzing the extent to which the consequences of the monetary policy response to a change in oil prices 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, defines the model space and studies its basic properties. Section 5 reports the results of the policy evaluation exercise. Section 6 concludes.

2 Policy evaluation under model uncertainty

In this section, I summarize the techniques developed by BDW (2003, 2007) and BDNR to account for model uncertainty in the evaluation of alternative economic policies.[6] These are the techniques that will be employed in the exercise in section 5.

2.1 General Framework

The central idea of 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 lack of complete information about 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. Typically, this policy will be studied 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 by BDW (2003, 2007) entails computing the probability measure $\mu\left(\theta \mid d, p\right)$ from (2) by treating model uncertainty as any other form of uncertainty affecting θ. This can be done by eliminating 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 measure can be characterized as follows:

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

[6] This section only provides a brief explanation of the techniques that I will use in section 5 of the paper. For a more thorough description of these methods, see BDW (2003, 2007) and BDNR.

where $\mu(d \mid m)$ is the likelihood of the data given model m and $\mu(m)$ 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 $l(\theta)$. The previous discussion implies that the measure incorporating model uncertainty into the analysis is:

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

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

In addition to the model averaging approach, BDW (2007) propose additional ways of communicating information about the effects of different policies in an environment characterized by model uncertainty. The introduction of these additional statistics is motivated by several considerations. First, the policymaker might want to investigate 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 across the different specifications composing the model space.

For all of these reasons, it could be useful to enrich the policy evaluation exercise by including additional measures that are able to offer a broader picture of the effects of a policy under alternative representations of the economy. 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 on the

[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.

robustness of the selected policy rule over different models. Action dispersion, on the other hand, measures how the optimal policy differs across alternative models. A distinct optimal policy can be computed for any given model, so that a range of different policies can be obtained from the elements of a model 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 the minimax and minimax regret criteria. These approaches are based on the idea that policymakers might be interested in obtaining information about policy rules that are not optimal, but that work well in some other directions or aspects of the policy analysis. In particular, these criteria address a concern for controlling the maximum losses that can be incurred under alternative policies in an environment characterized by model uncertainty.

The minimax approach has been largely used by Hansen and Sargent (2001a, 2001b, 2008) as the basis for robustness analysis in macroeconomics. In the policy evaluation exercise performed in section 5, I will follow BDNR and define the minimax policy choice as the one solving:

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

Because it always assumes the worst possible scenario in assessing alternative policies, the minimax criteria has been criticized for being extremely conservative. To avoid this issue, the literature has introduced the concept of minimax regret, which is based on the relative (rather than absolute) loss associated with a given policy. Following again BDNR, the minimax regret policy rule will be obtained as the solution to the following problem:

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

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 exposition of the properties of the minimax and minimax regret criteria and describe some applications 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) focus on the demand-side effects of changes in oil prices. In these frameworks, the direct and immediate consequence 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 assumption of rigid wages. Thus, wage rigidity is the main channel through which oil price variations affect output in these models. A similar explanation has been proposed by Blanchard and Gali (2007), which assume price rigidities in addition to 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. Other contributions 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 originate from 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. A rise in oil prices will decrease demand for some goods, but possibly increase it 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.

[8] Extensive reviews of the different mechanisms that have been proposed to explain the impact of oil prices on the economy are provided by Mork (1994), Hamilton (2005), Segal (2007) and Kilian (2008).

4 Model Uncertainty

I consider the problem of a policymaker who wants to investigate possible policy responses to changes in the price of oil. 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 relevant 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 through the assumption of 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 data generating process, the policymaker considers three different approximating frameworks that incorporate the main features of each one of these representations of the economy. These frameworks are in the spirit of the empirical literature on monetary policy, along the lines of King, Stock and Watson (1995), Rudebusch and Svensson (1999), Cogley and Sargent (2005) and Primiceri (2006). Each framework consists of two equations, one for the output gap and one for the inflation rate, and includes the following variables: the output gap (y_t), core CPI inflation (π_t), the interest rate (i_t), which is the policy instrument, and real oil price changes (s_t).[9]

The S approximating model is described by the following equations:

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

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

[9] The use of core CPI inflation follows Blanchard and Gali (2007) and Clark and Terry (2010).

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 standard setup in which nominal wages are set in advance (a thorough description of this type of setup can be found in Woodford, 2003).

The BG approximating model is described by the following equations:

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

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

This is a new Keynesian type of framework, similar to the model used in Rudebusch and Svensson (1999) under the assumption of backward expectations, with the only difference being the inclusion of real oil price changes. The specific form of equation (12) enforces that the sum of the coefficients on core CPI inflation is equal to one. This assumption is common in the new Keynesian literature, and follows from the theory on a vertical long run Phillips curve.

Finally, the H approximating model is described by the following equations:

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

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

This econometric framework reflects the theory on the independence of real variables from money and inflation, and 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, for an exhaustive treatment). Equations similar to (13) have been frequently used by Hamilton (1983, 2003, 2005) to estimate the effects of oil prices on output.

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

4.1 The process for the real price of oil

Equations (9) - (14) include the variable s_t, real oil price changes, defined as:

$$s_t = p_{s,t} - p_{s,t-1}$$

where $p_{s,t}$ is the level of the real price of oil at time t. I assume that this variable follows the exogenous AR(1) process:

$$p_{s,t} = \delta_t + \rho p_{s,t-1} + o_t \tag{15}$$

in which the intercept δ_t is allowed to change over time according to:

$$\delta_t = \delta_{t-1} + \varepsilon_{\delta,t} \tag{16}$$

The shocks o_t and $\varepsilon_{\delta,t}$ are assumed to be uncorrelated over time and with each other, and to have zero mean and variances σ_o^2 and $\sigma_{\varepsilon_\delta}^2$ respectively. Notice that while o_t represents a transitory oil price shock, $\varepsilon_{\delta,t}$ permanently affects the level of the real price of oil.

The mean shifting representation described by $(15)-(16)$ aims to capture the nonlinearities that seem to characterize the behavior of the real price of oil.[10] This process is similar to the one adopted by Kim and Loungani (1992), in which the intercept δ is constant but the shock o_t is allowed to be correlated over time, and is a generalization of the representation used in Blanchard and Gali (2007), which simply set $\delta_t = 0$ for any time t.

From the definition of the process for $p_{s,t}$, we have that:

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

where $\xi_t = (\varepsilon_{\delta,t} + o_t - o_{t-1})$ has zero mean and variance $\sigma_\xi^2 = \left(2\sigma_o^2 + \sigma_{\varepsilon_\delta}^2\right)$. The parameters of the representation in $(15)-(16)$ can be jointly estimated using MCMC methods. For the policy evaluation exercise conducted in this work, the policymaker only needs to know the values of σ_o^2, $\sigma_{\varepsilon_\delta}^2$ and ρ, while the entire history of δ_t is not necessary. I estimated the process in $(15)-(16)$ using U.S. postwar data in a previous paper (Rondina, 2010). Therefore, in the empirical analysis I will simply use the values obtained therein; the reader is encouraged to refer to Rondina (2010) for a detailed description of the estimation procedure.

In this paper, I follow Kim and Loungani (1992) and Blanchard and Gali (2007) and assume that the real price of oil follows an exogenous process. This assumption might seem quite restrictive, especially in regards to the changes in oil prices that happened during the last decade. However, the specific process described by $(15)-(16)$ allows for a certain degree of flexibility, which reduces the limits imposed by the assumption of exogeneity. See Rondina (2010) for a more extensive discussion of this issue.

4.2 The policy rule

I assume that the policymaker employs a simple nominal interest rate rule in the form:

$$i_t = g_\pi \pi_t + g_y y_t + g_i i_{t-1} + g_s s_t \tag{18}$$

[10] See Pindyck (1999) for a discussion. Kim and Loungani (1994) and Blanchard and Gali (2007) also suggest that the real price of oil would be better described by a process that can accommodate nonstationary.

Simple rules in the form of (18) are used in BDW (2007) and in other contributions in the literature on monetary policy as, for instance, Levin, Wieland and Williams, (1998). The main difference with previous studies adopting similar rules is the addition of the last term, which allows policymakers to respond to changes in the real price of oil.

Following standard assumptions in the monetary rules literature, policymakers' choice of the parameters g_π, g_y, g_i and g_s in (18) affect their welfare defined in terms of an expected loss function. I follow the literature on policy evaluation under model uncertainty (see for instance BDW, 2003 and 2007, and Cogley et al., 2009) and I assume that policymakers' losses are determined by a weighted sum of the unconditional variances of the variables of interest:

$$R = var\left(\pi_\infty\right) + \lambda_y var\left(y_\infty\right) + \lambda_i var\left(\Delta i_\infty\right) \tag{19}$$

The parameters λ_y and λ_i in (19) represent the weights attached to the volatility of the output gap and of the changes in the policy instrument relative to the volatility of core CPI inflation. I will assume that $\lambda_y = 1$ and $\lambda_i = 0.1$ as in BDW (2007) and Cogley et al. (2009); this choice is consistent with the literature using similar loss functions, see for instance Levin and Williams (2003). The last term in (19), which accounts for interest rate variations, is commonly introduced with a lower weight compared to the other variables, and its role is to avoid extreme changes in the nominal interest rate.

In the policy evaluation exercise in section 5, different values of R will be calculated based on alternative conditioning assumptions made via specification of a policy and/or a model.

4.3 The model space

The space of all model specifications included into the analysis is defined based on different forms of model uncertainty that policymakers view as relevant in the environment under consideration.

- Theory uncertainty. The first, and most important, form of model uncertainty the monetary authority is concerned about is theory uncertainty. Theory uncertainty refers to the 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 originate three different classes of models that span 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 decision process even if the true model of the economy was known.

In more detail, I follow BDW (2007) and 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 incorporates this form of uncertainty by estimating equations (9), with one, two, three and four lags of y, unanticipated inflation and s, (11) with one, two, three and four lags of y, the real interest rate and s, and (13) with one, two, three and four lags of y and s. In the same way, he estimates equations (10) and (14) with one, two, three and four lags of y, π and i, and (12) with one, two, three and four lags of y and π only, while in all models the inflation equation will be allowed to include zero, one, two, three and four lags of s. The number of lags of s_t used in the estimation procedure reflects the fact that all the theories under consideration postulate an effect of oil prices on output, while the impact on core CPI inflation, which excludes energy prices, is not obvious.[11] Table 1 summarizes the number of lags included in the estimation of each approximating model.[12]

Table 1 - Lags used in the estimation procedure

	$l.h.v$	y	π	i	s
S	y	1-4	1-4	0	1-4
	π	1-4	1-4	1-4	0-4
BG	y	1-4	0	1-4	1-4
	π	1-4	1-4	0	0-4
H	y	1-4	0	0	1-4
	π	1-4	1-4	1-4	0-4

Notes: 1. Number of lags of the right-hand variables included in each equation (with left-hand variable y or π) for each approximating model: Solow (S), Blanchard-Gali (BG) and Hamilton (H).
2. In the S model, the right-hand variables are output gap (y), unanticipated inflation (denoted as π here for simplicity of exposition) and real oil price changes (s) in the output equation and output gap (y), inflation (π), interest rate (i) and real oil price changes (s) in the inflation equation. In the BG model, the right-hand variables are output gap (y), real interest rate (denoted as i here for simplicity of exposition) and real oil price changes (s) in the output equation and output gap (y), inflation (π) and real oil price changes (s) in the inflation equation. Finally, in the H model, the right-hand variables are output gap (y) and real oil price changes (s) in the output equation and output gap (y), inflation (π), interest rate (i) and real oil price changes (s) in the inflation equation.

[11] In Blanchard and Gali (2007), for instance, real oil prices do not affect core CPI inflation if the assumption of real wage rigidities is dropped, even if Calvo pricing is still adopted. Solow (1980) also acknowledges the possibility that oil price shocks do not change core inflation (but he says that this situation is very unlikely). Finally, Kim and Loungani (1992) focus on the effects of oil prices on the production side of the economy, while their impact on the price level is not discussed.

[12] Notice that the econometric models considered in this work are similar to those used in Rondina (2010) which studies the history of postwar US policy decisions under the assumption of model uncertainty and learning. However, specification uncertainty is not considered in this previous work.

4.4 Basic properties of the model space

The different forms of uncertainty that the policymaker decides to incorporate in the analysis originate a model space M composed of $30,720$ models, $5,120$ for M^{BG} and M^{H} and $20,480$ for M^{S}. Because the equations of the Solow model include more variables relative to the Blanchard-Gali and Hamilton models, the Solow class of models is four times larger than the other two classes. This requires taking a stance on the definition of the prior probabilities to be used in computing the models' posteriors, as I will discuss later in this section.

I estimated all the econometric specifications in M using ordinary least squares. The data employed in the estimation goes from $1973:I$ to $2008:II$, with observations from $1971:I$ to $1969:IV$ used to provide lags. While the available data covers a larger time period, the choice of the sample to be used in the policy evaluation exercise was motivated by the fact that in this framework all parameters of each approximating model (with the exception of the intercept in the process for the real price of oil) are assumed to be time invariant. This implies that the necessity to have a long enough sample for the estimates to be meaningful must be balanced with the fact that over longer time periods the economy is more likely to have undergone relevant structural changes, which could have been reflected into variations in the model parameters. More specifically, the starting date was motivated by the fact that until the early $1970s$ oil prices were subject to strong price controls, which were likely affecting the behavior of oil users. The ending date was selected so that the recent financial crisis, and the near zero interest rate policy adopted by the Federal Reserve in response to it, would not affect the estimations. A detailed description of the data used in the empirical analysis is given in Appendix 1.

In estimating the S and BG specifications, I assumed that the public forms expectations using a backward-looking approach, so that: $E_{t-1}(\pi_t) = \frac{1}{4}\sum_{j=1}^{4}\pi_{t-j}$. This assumption might seem restrictive, but in an environment characterized by model uncertainty forward looking expectations rise a number of questions related to whether and to what extent the private sector should share policymakers' uncertainty about the true data generating process. Again, Rondina (2010) discusses this issue in more detail.

The posterior probabilities were computed according to (4), using the estimated models and the approximation suggested by Raftery (1995). In this way, the obtained posteriors are proportional to the models' BIC adjusted likelihoods, with a factor of proportionality that is equal to their respective prior probabilities. In the baseline exercise, I assumed that policymakers' main focus is on theory uncertainty, so that a prior probability of 1/3 was attached to each class of models regardless of the number of specifications that each class contains. Therefore, since the Solow class of models includes a much larger number of specifications, each one of them received a considerably lower prior relative to the specifications in the other classes. In more detail, a prior of 1/15360 was attached to each element of M^{BG} and M^{H}, while a prior of 1/61440 was attached to each model in M^{S}. In the empirical section of the paper, I will study the robustness of the results to an alternative definition of these prior probabilities.

Figure 1 - Posterior probabilities

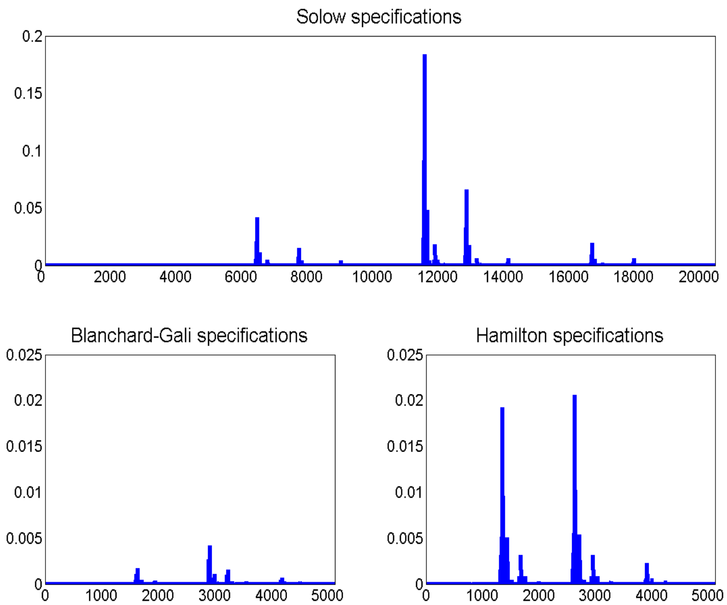

Notes: 1. Posterior probabilities for each model specification in M. The top panel reports the Solow class of models, and the bottom panels the Blanchard-Gali and Hamilton classes. Notice that, for a better exposition of the results, the scale in the bottom panels is different from the top panel.

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 each class of models) is reported in Table 2.

3. Model numbers are explained in Appendix 1.

Table 2 - Sum of posterior probabilities in each class of models

	M^S	M^{BG}	M^H
Sum of posterior probabilities	0.8585	0.0301	0.1114
Sum of prior probabilities	1/3	1/3	1/3
No models	20,480	5,120	5,120

Note: Sum of posterior probabilities in each class of models. The posterior probabilities have been computed from (4) using the approximation proposed by Raftery (1995) and the priors described in the main text. The posteriors have been rescaled so that they add up to one across the model space M.

Figure 1 reports the posterior probabilities for each model specification in M. Table 2 provides information on the sum of these posteriors in each class of models. Posterior probabilities have been rescaled so that they add up to 1 across the model space. Thus, table 2 can be interpreted as the probability that the true data generating process follows the Solow, Blanchard-Gali or Hamilton theory on the predominant channels through which oil prices are assumed to affect the economy. In this sense, it is clear that the data favors the Solow theory, since this class of models incorporates 85.85% of the posterior probability. However, the posteriors attached to the Hamilton and Blanchard-Gali theories, while considerably lower relative to the Solow class, are still largely different from zero. In addition, figure 1 also shows that a few specifications, belonging to different classes of models, exhibit posterior probabilities that are actually comparable with each other. For these reasons, a policymaker concerned about model uncertainty should not discard any of these theory as the possible true representation of the economy, but rather look for a policy rule that is able to perform relatively well in all of them.

From figure 1, it is evident that each class of models is characterized by an handful of specifications that have higher posterior probabilities, and a large number of them that, on the contrary, have near zero posteriors. Given the large number of models in M, policymakers might want to restrict the model space and focus only on those specifications that offer a plausible representation of the economy. In this choice, decision makers face a tradeoff between allowing for a sufficiently large degree of model uncertainty, and making the policy evaluation exercise cumbersome and possibly even not informative.[13] Here, I follow BDW (2007) in the procedure used to restrict the analysis to a smaller model space.[14] This procedure entails computing the relative posterior of a model within a class, defined as:

$$P_m = \frac{\mu(m \mid d)}{\sum\limits_{m \in C} \mu(m \mid d)} = \frac{\widehat{L}_m}{\sum\limits_{m \in C} \widehat{L}_m} \tag{20}$$

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. The second equality follows from the fact that in this setup posterior probabilities are proportional to BIC-adjusted likelihoods and that, within each class of models C, all models have the same prior. In words, this formula rescales the posterior probabilities so that they add up to one within each class of models. The measure obtained from (20) is then used to identify the models that have the highest relative posterior probabilities within each class. In this work, these models will be defined as those for which P_m is at least $1/100 = 1\%$ of the model with the highest P_m in the class. The policy evaluation

[13]Many of the specifications with near zero posteriors are very unstable, and exhibit infinite losses under a wide range of policies. For this reason, they might dominate the policy evaluation exercise, despite the fact that their posterior probability is essentially zero.

[14]This approach is based on the "Occam's window" technique originally proposed by Madigan and Raftery (1994).

exercise developed in the next section will focus on this subset of model specifications.[15]

Table 3 - Relative posterior probability P_m

	M^S	M^{BG}	M^H
(1) Minimum P_m	2.02×10^{-24}	7.34×10^{-26}	2.47×10^{-22}
(2) Q1 P_m	1.53×10^{-15}	2.16×10^{-18}	3.89×10^{-14}
(3) Median P_m	6.40×10^{-12}	9.08×10^{-12}	1.74×10^{-10}
(4) Q3 P_m	3.92×10^{-9}	8.58×10^{-7}	8.59×10^{-8}
(5) Maximum P_m	0.2137	0.1378	0.1849
(6) No. models with $P_m >$ ($\texttt{max}$ P_m)/100	55	87	56
(7) Sum of P_m models with $P_m >$ ($\texttt{max}$ P_m)/100	0.8581	0.8827	0.9274
(8) Sum of P_m for models in top quartile	1.0000	0.9998	1.0000
(9) Sum of P_m for models in bottom 3 quartiles	3.25×10^{-6}	1.97×10^{-4}	1.89×10^{-5}
(10) Sum of P_m	1	1	1
(11) No. models	20,480	5,120	5,120

Note: The relative posterior probability P_m is defined by (20). The sum of P_m for each class of models equals one by construction.

Table 3 provides some summary statistics on the distribution of the relative posterior probabilities for each class of models. This table clearly shows that, in each class, a restricted number of specifications cover almost the entire posterior probability for the class. Indeed, the first three quartiles only contain specifications with relative posteriors that are essentially zero, while the sum of P_m for the first quartile is nearly one in all classes. The number of specifications for which P_m is at least 1% of the model with the highest P_m, reported in line (6), is very small relative to the size of each class, but these few specifications still cover a very high relative posterior, as shown in line (7). For the policy evaluation exercise in the next section, the model space M and the classes of models M^S, M^{BG} and M^H are redefined to incorporate only the models with the highest relative posterior probability. Therefore, the new model space includes 198 specifications, while M^S, M^{BG} and M^H are composed of 55, 87 and 56 models respectively. A more detailed description of the model specifications used in the policy evaluation exercise, and the definition of the new model space and classes of models are provided in Appendix 1.

[15] The factor that is used in BDW (2007) to define the set of models with high posterior probability is 1/20. The reason why I set a lower threshold is that, in this context, a large number of models, with posterior probability different from zero as a group, do not get captured by the 1/20 threshold. Since I will use the subset of models with high posterior probabilities for the policy evaluation exercise in the next section, the lower threshold of 1/100 allows me to have a group of models that provide a better representation of the original model space M.

Table 4 - Parameter estimates for the models with the highest posterior probability

(A) Output equation

	α_{y_1}	α_{y_2}	α_{y_3}	α_{π_1}	α_{π_2}	α_{i_1}	α_{i_2}	α_{s_1}	$\overline{R}^2$	DW	$s.e.$
S model	1.183 (0.007)	−0.024 (0.016)	−0.221 (0.006)	0.030 (0.002)	−0.162 (0.002)	$n.a.$	$n.a.$	−0.0014 (0.000)	0.90	1.92	0.49
BG model	1.132 (0.006)	−0.028 (0.014)	−0.206 (0.006)	$n.a.$	$n.a.$	0.088 (0.002)	−0.152 (0.002)	−0.0014 (0.000)	0.90	1.98	0.50
H model	1.187 (0.007)	−0.107 (0.017)	−0.186 (0.007)	$n.a.$	$n.a.$	$n.a.$	$n.a.$	−0.0011 (0.000)	0.89	2.03	0.54

(B) Inflation equation

	β_{y_1}	β_{y_2}	β_{π_1}	β_{π_2}	β_{π_3}	$\beta_{\pi 4}$	β_{i_1}	β_{i_2}	β_{i_3}	β_{i_4}	$\overline{R}^2$	DW	$s.e.$
S model	0.137 (0.004)	$n.a.$	0.290 (0.008)	0.073 (0.008)	0.349 (0.007)	0.235 (0.008)	0.438 (0.009)	−0.461 (0.014)	0.379 (0.015)	−0.398 (0.010)	0.78	1.80	1.83
BG model	0.622 (0.029)	−0.392 (0.027)	−0.637 (0.006)	−0.590 (0.006)	$n.a.$	$n.a.$	$n.a.$	$n.a.$	$n.a.$	$n.a.$	0.39	2.01	2.21
H model	0.137 (0.004)	$n.a.$	0.290 (0.008)	0.073 (0.008)	0.349 (0.007)	0.235 (0.008)	0.438 (0.009)	−0.461 (0.014)	0.379 (0.015)	−0.398 (0.010)	0.78	1.80	1.83

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, but are not reported for clarity of exposition.

2. In Panel (A), output gap is the dependent variable, α_{yj} is the coefficient on output gap at lag j, α_{ij} and $\alpha_{\pi j}$ are the lag j coefficients on the annual real interest rate and unanticipated inflation respectively, and α_{sj} is the coefficient on real oil price changes at lag j. In Panel (B), for the S and H specifications, 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, and β_{ij} is the coefficients on the annual nominal interest rate at lag j. For the BG specification, the change in inflation is the dependent variable, β_{yj} is the coefficient on output gap at lag j and $\beta_{\pi j}$ is the coefficient on the change in inflation at lag j.

3. The sample is composed of quarterly data from 1973:I to 2008:II, for a total of 142 observations. Inflation is the annualized change in core CPI; the output gap is the difference between real GDP and the CBO estimate of potential GDP, both in lags; the interest rate is the average annual Federal funds rate; real oil price changes are the annualized change in the real price of oil, computed as 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.

Finally, table 4 reports the estimated coefficients for the specification with the highest posterior probability in each class of models. As I mentioned before, posterior probabilities are proportional to model specific BIC-adjusted likelihoods. It follows that the specifications presented in table 4 correspond to those that would have been selected within each class using BIC as the selection criterion. Notice that in these specifications real oil price changes enter in the output equation with only one lag, and they do not enter in the inflation equation. However, the subspace of models with high posterior probabilities used in the policy analysis includes specifications with a higher number of lags of the oil measure in both equations. Again, see Appendix 1 for further details on the elements of the restricted model space.

4.5 Simple rules

This work aims to compare the performance of alternative policy rules in an environment characterized by uncertainty on the way oil prices affect economic variables. Thus, after having described the space of models under consideration, the second step is defining the set of policies to be evaluated. As previously mentioned, I assume that policymakers only consider simple policy rules in the form of (18).

The first rule included in the set of policies under analysis is the one originally proposed by Taylor (1993) (from now on denoted as OT rule):

$$i_t = 1.5\pi_t + 0.5y_t \tag{21}$$

This policy rule is widely used in the literature and was likely also implemented in practice, so it will be considered as a benchmark. In addition to the OT rule, policymakers might want to study the performance of policies that are to some extent optimal under the theories they regard as possibly generating the data. To obtain these policy rules, I followed BDW (2007) and used the specification with the highest posterior probability in each class of models. More specifically, I computed these rules by performing a grid search of the parameters g_π, g_y, g_i and g_s in (18) that minimize the conditional expected loss:

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

for each of the three models described in table 4. I restricted this search to rules in which the long run effect of output and core CPI inflation on the nominal interest rate is the same as in the Taylor rule.[16] No restrictions were imposed on the coefficient on real oil price changes, g_s. In other words, I assumed that the monetary authority wants to evaluate the performance

[16] More specifically, I performed a grid search only on values of g_π, g_y and g_i that satisfy: $g_\pi/(1-g_i) = 1.5$ and $g_y/(1-g_i) = 0.5$.

of the Taylor rule relative to alternative simple rules which differ from the original Taylor rule only in terms of interest rate smoothing and the (possible) response to oil prices. This exercise provides a clear picture of the impact that reacting to changes in the real price of oil has on policymakers' losses, and seems to be the most appropriate in a context characterized by uncertain on the way in which oil prices affect economic variables.[17]

The simple policy rules obtained from the described procedure, denoted as S rule, BG rule and H rule, are reported in table 5.

Table 5 - Policy space: the simple policy rules

	S rule	BG rule	H rule
g_y	0.1995	0.4635	0.2670
g_π	0.5985	1.3905	0.8010
g_i	0.6010	0.0730	0.4660
g_s	-0.0071	-0.0194	-0.0040
Exp. loss	31.688	22.706	19.771
Long run			
$\widetilde{g}_\pi$	1.5	1.5	1.5
$\widetilde{g}_y$	0.5	0.5	0.5
$\widetilde{g}_s$	-0.0178	-0.0209	-0.0075

Notes: 1. Simple rules in the form described by (18). These rules were obtained by grid search of the coefficients in (18) that minimize (22) under the restrictions $g_\pi/(1-g_i) = 1.5$ and $g_y/(1-g_i) = 0.5$ for the specification with the highest posterior probability in each class of models.
2. The long run effect of y, π and s on the nominal interest rate is defined as: $\widetilde{g}_k = g_k/(1-g_i)$, $k = y, \pi, s$.

The simple rules reported in table 5 offer some relevant insights on the differences in the optimal policy response to oil prices in each of the three theories under consideration. In particular, we can compare the short run and long run effects of oil prices on the nominal interest rate that these three policies imply. As expected, the BG rule recommends the strongest response to changes in the real price of oil, both in the short run and in the long run. Indeed, the Blanchard-Gali theory assumes that the economy is characterized by a number of rigidities that

[17] In a previous version of the paper, I was comparing the original Taylor rule to the optimal simple rules obtained by minimizing (22) with no restrictions on the values of the coefficients g_y and g_π. However, I found that exercise to be less informative than the one performed here. Indeed, the differences in performance between the alternative simple rules and the original Taylor rule were largely driven by their different response to output and inflation, and it was difficult to discern the role of the reaction to changes in the real price of oil. Here, this is not the case, because the long run response to the output gap and core CPI inflation is set to be equal in all the rules considered in the policy evaluation exercise.

have the potential to amplify the impact of oil prices on the variables of interest to policymakers. At the same time, in this theory the policy instrument can affect the output gap directly so that, by responding to changes in the real price of oil, policymakers are able to contrast the effects of this variable on the real economy. On the other hand, in the Solow theory monetary policy has an impact on the output gap only indirectly through unanticipated inflation. Nonetheless, the policy suggested by this theory still implies a relatively large reaction of the nominal interest rate to oil prices, especially in the long run. Finally, in the Hamilton theory policymakers cannot modify the output gap with their policy choices. For this reason, the interest rate response to the oil variable is much smaller, and directed to contrast its effects on core CPI inflation only.

Figure 2 - Impulse responses: simple rules and no action

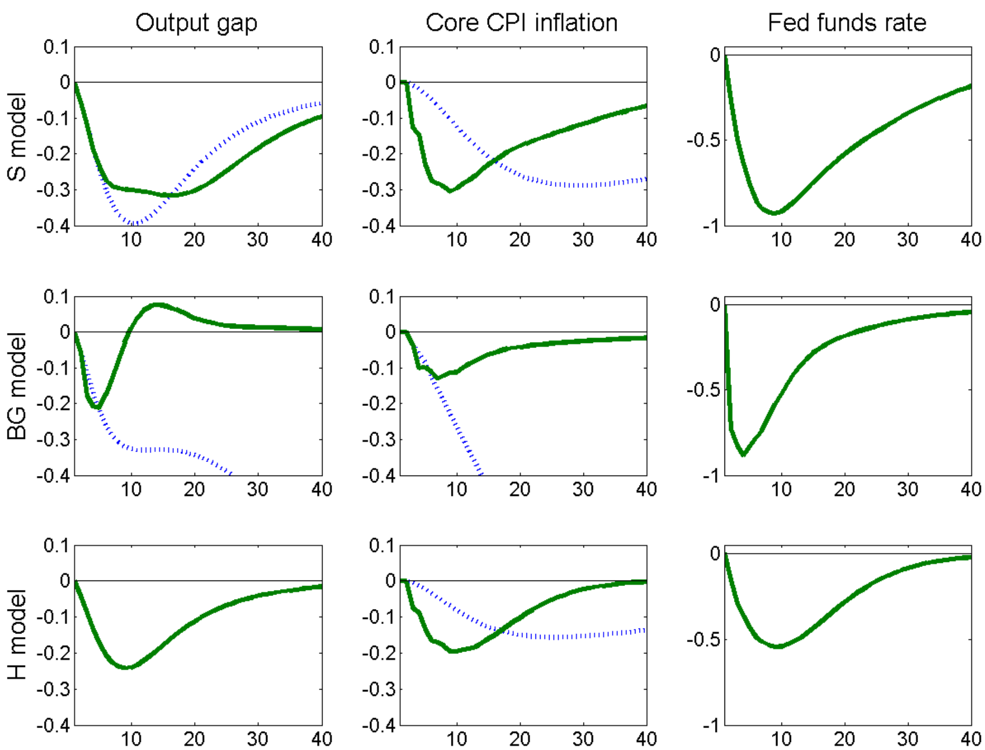

Note: Response of the output gap, core CPI inflation and the Federal funds rate to a 10% increase in the real price of oil. The first columns reports the output gap, the second column core CPI inflation and the last column the Federal funds rate. Each row represents a different model and relative policy 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 (18) are all set equal to zero (dashed line).

For a better understanding of the implications of the simple rules described in table 5, I studied the policy response to a 10% increase in the real price of oil that each of them recommends. More specifically, I investigated the response of output, core CPI inflation and the Federal Funds rate in each of the models described in table 4, when the policymaker implements the respective optimal simple rule. In each model, the impact of the policy response to the change in oil prices 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 (18) 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 shock is model dependent. The results of this exercise are reported in figure 2; some further analysis of the policy responses implied by each of the rules described in table 5 is provided in Appendix 2.

A few things can be observed from figure 2. First, as discussed the recommended response to oil prices is stronger in the Solow and Blanchard-Gali theories relative to the Hamilton theory. Second, in the Blanchard-Gali model if policymakers do not respond to the change in the real price of oil, both the output gap and core CPI inflation quickly diverge towards infinite negative values. Thus, in this model policymakers must react to changes in oil prices to preserve the stability of the variables of interest. Third, figure 2 shows that the ability of policymakers to contrast the effects of a change in the real price of oil on the output gap is quite different depending on the model of the economy under consideration. For this reason, the exercise reported in this figure provides some additional insights on the debate between Bernanke et al. (1997, 2004) and Hamilton and Herrera (2004) over the role of monetary policy in the declines in output that followed most of the oil price shocks of the postwar period. While Bernanke et al. (1997, 2004) suggest 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) argue that output would have decreased no matter what policy had been implemented. Figure 2 reports an impulse response exercise that is very similar to those studied by Bernanke et al. (1997, 2004) and Hamilton and Herrera (2004), and the panels in the first column of this figure are actually consistent with the results of these contributions. In more detail, if the true model of the economy is the BG model, then figure 2 shows that policymakers can successfully reduce the downfall in output caused by an oil price shock by implementing an expansionary policy rule. This conclusion supports the position of Bernanke et al. (1997, 2004). On the other hand, if the true data generating process is either the H model or the S model, then policymakers are not able to avoid the decrease in output caused by a change in oil prices, and a more expansionary policy rule brings no benefits to the real economy, which is the opinion expressed by Hamilton and Herrera (2004). Thus, this exercise provides evidence that both positions can be correct, depending on which theory is regarded as the one generating the data.

The simple rules reported in table 5 have been selected to minimize losses in a specific model belonging to one of the three theories under consideration. However, their performance in the other specifications included in their same class or in the other classes of models is not obvious, and policymakers might be interested in evaluating whether the adoption of one of them offers advantages relative to the implementation of the OT rule. This exercise is carried out in the next section, using the measures that have been previously described in section 2.

5 Policy Evaluation

In a context in which the monetary authority 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? These questions will be investigated in this section.

A large part of the policy evaluation exercise performed in this section is based on the study of expected losses conditional on a given model specification and policy rule, as defined in (22). In addition, the Bayesian portion of the analysis requires the computation of expected losses for the different classes of models and for the entire model space. As in BDW (2007) and Cogley et al. (2009), these will be obtained by taking a weighted average of the model specific conditional expected losses, using posterior probabilities as weights. Thus, the expected loss across the entire model space when model uncertainty is incorporated into the analysis will be defined as:

$$\widehat{R} = \sum_{m \in M} \widehat{R}_m \mu \left(m \mid d \right) \tag{23}$$

Using the same approach, the expected loss for each class of models will be computed as:

$$\widehat{R}_C = \sum_{m \in C} \widehat{R}_m \mu \left(m \mid d \right) = \frac{\sum_{m \in C} \widehat{R}_m \widehat{L}_m}{\sum_{m \in C} \widehat{L}_m} \tag{24}$$

where again $\widehat{L}_m$ is the BIC adjusted likelihood for model m, and the second equality follows from the fact that within each class of models all specifications have the same prior probability.

5.1 Outcome dispersion

Outcome dispersion measures the variation in loss that occurs when considering the effects of the same policy rule in different model specifications. Table 6 reports the properties of the distribution of losses for each class of models under each of the four policy rules included into the analysis (OT rule, S rule, BG rule and H rule). Table 7 provides a description of the same distribution across the whole model space. Finally, figure 3 offers a visual representation of the information presented in these two tables.

Table 6 - Distribution of model losses under each of the policy rules

Class of models	M^S				M^{BG}				M^H			
Policy rule	OT	S	BG	H	OT	S	BG	H	OT	S	BG	H
(1) Mean	42.92	39.48	40.88	40.31	68.46	46.87	40.67	57.29	32.90	30.86	32.49	30.54
(2) St. deviation	11.45	10.72	9.97	11.34	36.23	23.31	20.64	29.80	11.51	10.31	9.20	11.16
(3) Minimum	26.21	25.54	28.57	24.59	27.42	22.60	21.20	26.22	19.07	19.90	22.22	18.15
(4) Q1	34.30	31.00	33.37	31.54	40.90	27.90	24.92	34.22	25.11	22.92	25.66	22.21
(5) Median	41.26	36.62	38.51	38.10	49.46	35.34	33.01	43.75	30.08	28.31	29.05	28.39
(6) Q3	49.40	44.92	44.74	47.10	96.23	62.15	52.92	80.40	37.78	34.36	35.23	34.56
(7) Maximum	78.55	75.67	75.53	77.02	157.41	127.63	130.24	132.26	67.20	64.32	63.72	65.59
(8) P. w. average	38.70	35.31	36.97	36.07	57.56	36.80	31.87	46.43	26.85	25.30	27.19	24.71
(9) N. of models	55				87				56			

Notes: 1. Distribution of model specific losses for each class of models under the Taylor rule and the three simple rules described in table 5: the S rule, the BG rule and the H rule.

2. Rows (1) - (7) report basic statistics of the distribution of losses for each class of models under each policy rule. Row (8) reports the posterior weighted average loss, computed using (24).

3. The composition of each class of models is described in Appendix 1.

In the environment under analysis, the simple rules described in table 5 perform better than the OT rule in terms of the first and second moments of the distribution of losses that they generate. In particular, table 6 shows that while the OT rule implies higher and more disperse expected losses in all classes of models, its performance is significantly worse than the other rules in the BG class. Among the three simple rules, the H rule delivers a higher mean and standard deviation of losses than the S and BG rules in the BG class, while all of them imply similar losses in the other two classes. The considerably lower standard deviation of expected losses that can be attained by adopting the S or BG rule should be a characteristic of particular interest to policymakers in an environment characterized by uncertainty on the model that generates the data.

Table 7 - Distribution of losses across the model space

	OT rule	S rule	BG rule	H rule
(1) Mean	51.308	40.287	38.414	45.009
(2) Standard deviation	29.866	18.521	15.840	24.289
(3) Minimum	19.068	19.900	21.198	18.152
(4) Q1	32.592	27.894	26.692	29.511
(5) Median	41.377	33.323	33.948	36.224
(6) Q3	55.738	48.493	44.579	51.304
(7) Maximum	157.413	127.634	130.243	132.259
(8) Posterior weighted average	37.866	34.164	35.647	35.030
(9) N. of models	198	198	198	198

Notes: 1. Distribution of model specific losses across the entire model space under the Taylor rule, the S rule, the BG rule and the H rule.
2. Rows (1) - (7) report basic statistics of the distribution of losses under each policy rule. Row (8) reports the posterior weighted average loss, computed using (23).
3. The composition of the model space is described in Appendix 1.

Figure 3 - Outcome dispersion for each policy rule

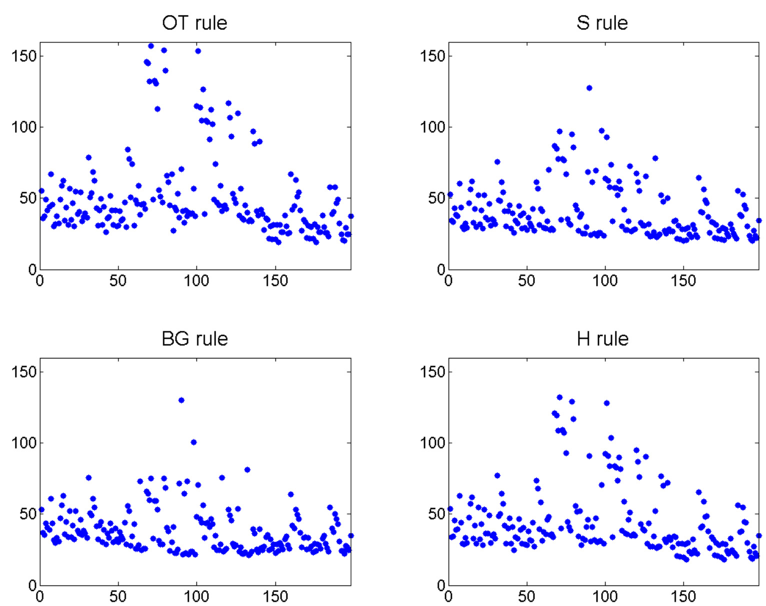

Notes: 1. Model specific expected losses under the original Taylor (OT) rule, defined in (21), and the S, BG and H rules described in table 5.
2. The summary statistics for the distribution of losses in each class of models are reported in table 6. The summary statistics for the distribution of losses across the model space are reported in table 7.
3. Models from 1 to 55 belong to the S class, from 56 to 142 to the BG class, and from 143 to 198 to the H class. Additional information on the model numbers is provided in Appendix 1.

In terms of posterior weighted average losses, the differences in performance between rules are reduced when we consider the entire model space because the BG class of models, in which the disparities are larger, covers a smaller posterior probability relative to the other two classes. Despite this, the posterior weighted average loss delivered by the OT rule is the largest among the policy rules under consideration. Notice that this measure is the only one in tables 6 and 7 that is computed using the models' posterior probabilities, while all the other information is obtained by assigning an equal weight to all specifications in the model space. The posterior weighted average loss is the value that is naturally used for policy evaluation in the Bayesian approach. It follows that, in this environment, a Bayesian policymaker would select the S rule as the robust policy under model uncertainty.

Figure 4 - Model losses for each policy relative to the Taylor rule

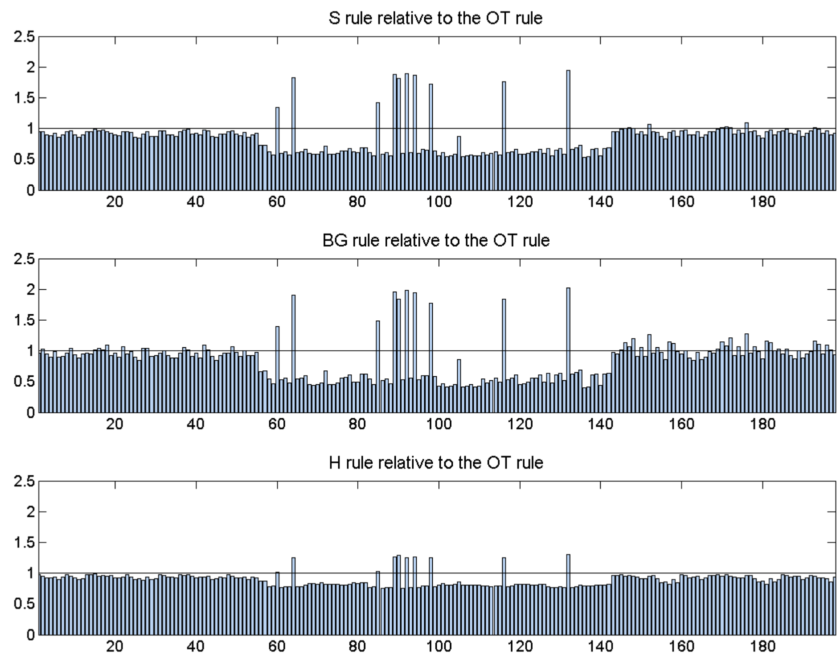

Notes: 1. Each panel reports the ratio between the loss generated by one of the simple policy rules described in table 5 and the loss generated by the original Taylor rule, for each specification in the model space.
2. Model numbers are explained in Appendix 1.

The performance of the OT rule relative to the alternative policy rules described in table 5 is further investigated in figure 4. For each model specification, this figure reports the ratio between the loss generated by the S, BG and H rules and the loss generated by the original Taylor rule. This exercise confirms that, in average, expected losses are lower under the alternative simple rules than under the OT rule, and that the largest improvement is attained in the BG class of models. At the same time, figure 4 also shows that there are a few specifications for which the OT rule implies lower losses compared to the S, BG and H policies. This is particularly true for some elements of the BG class, in which the original Taylor rule performs considerably better than the other policies, especially the S and BG rules. These are models in which the strong expansionary response to a change in oil prices suggested by the S and BG policies is considerably less effective in stabilizing the volatility of the variable of interest to policymakers than abstaining from any direct reaction to the change.

5.1.1 The role of inertia in the alternative policy rules

Tables $6-7$ and figures $3-4$ provide evidence of the fact that among the policy rules considered in the policy evaluation exercise, the OT rule is the one that delivers the highest mean and variance of the distribution of losses across the model space, and the highest posterior weighted average loss. Given this result, policymakers might be interested in investigating whether this difference in performance is due to the fact that the S, BG and H rules incorporate a response to changes in the real price of oil, while the OT rule does not.

The alternative simple rules considered in the policy evaluation exercise have been constructed so that the long run effect of the output gap and core CPI inflation on the nominal interest rate are the same as in the original Taylor rule. Therefore, the reduction in losses that these rules imply relative to the OT rule cannot be attributed to differences in the long run response to y_t or π_t. Nonetheless, in the S, BG and H policies this long run effect is attained through some degree of interest rate inertia, while this is not the case in the OT rule. Thus, the relevant question is whether the original Taylor rule performs worse than the alternative policy rules because it does not include 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 compared the results obtained in the first part of this section with the performance of a Taylor-type rule that adds 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:

$$i_t = g_i i_{t-1} + (1 - g_i)\widetilde{i}_t \qquad (25)$$

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

$$\widetilde{i}_t = \widetilde{g}_\pi \pi_t + \widetilde{g}_y y_t \qquad (26)$$

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

I focused on the distribution of losses induced by the policy rule defined in (25) and (26), with $g_i = 0.65$, $\widetilde{g}_\pi = 1.5$ and $\widetilde{g}_y = 0.5$. The value of g_i is in the range estimated by the empirical literature on inertial Taylor rules (see for instance Sack, 1998; Orphanides, 2001; Dueker and Rasche, 2004). As for the rules in table 5, I imposed the restriction that $\widetilde{g}_\pi$ and $\widetilde{g}_y$ are the same as in the original Taylor rule. The policy following from these assumptions, denoted as "inertial Taylor rule" (IT rule), is $i_t = 0.525\pi_t + 0.175y_t + 0.65i_{t-1}$. This rule is actually quite similar to the S rule, except for the lack of a response to changes in the real price of oil. For this reason, the comparison of the performance of the IT and S rules is particularly relevant for a better understanding of the impact that policymakers' reaction to oil prices has on outcome dispersion.

Table 8 - Distribution of losses across the model space

	OT rule	IT rule	S rule	BG rule	H rule
(1) Mean	51.308	54.801	40.287	38.414	45.009
(2) Standard deviation	29.866	36.418	18.521	15.840	24.289
(3) Minimum	19.068	17.477	19.900	21.198	18.152
(4) Q1	32.592	32.031	27.894	26.692	29.511
(5) Median	41.377	43.018	33.323	33.948	36.224
(6) Q3	55.738	59.073	48.493	44.579	51.304
(7) Maximum	157.413	184.666	127.634	130.243	132.259
(8) Posterior weighted average	37.866	37.409	34.164	35.647	35.030
(9) N. of models	198	198	198	198	198

Notes: 1. Distribution of model losses under different policy rules. The OT rule is defined by (21), the IT rule is defined by (25) and (26), with g_i= 0.65, $\widetilde{g}_\pi$= 1.5 and $\widetilde{g}_y$= 0.5, the S, BG and H rules are defined by (18), with coefficient values as reported in table 5.
2. Rows (1) - (7) report basic statistics of the distribution of losses under each policy rule. Row (8) reports the posterior weighted average loss, computed using (23).
3. The composition of the model space is described in Appendix 1.

Table 8 reproduces table 7 with the addition of the IT rule. In terms of loss dispersion, the IT rule performs even worse than the OT rule, although it still delivers a similar posterior weighted average loss. In comparison with the S rule, the IT rule generates a distribution of losses that has considerably higher mean, and almost twice the standard deviation. Thus, it is clear that the introduction of inertia in the Taylor rule is not able, per se, to improve the distribution of losses across the model specifications considered in the analysis. It follows that, in this environment, the direct response to changes in oil prices seems to play an important role in reducing the mean (simple and posterior weighted) and the volatility of expected losses for the elements of the model space.

5.1.2 The process for the real price of oil

The models described in $(9)-(14)$ incorporate oil prices in the form of the annualized change in the real price of oil. This variable is assumed to follow the process characterized by $(15)-(17)$. As already mentioned, in the baseline scenario expected losses were computed using values of ρ, σ_o^2 and $\sigma_{\varepsilon_\delta}^2$ that are consistent with the estimations performed in Rondina (2010). Alternative assumptions about the value of these parameters will affect policymakers' losses, and could alter the optimal policy response to oil prices recommended by the different specifications included in the model space. These changes have the potential to modify the previous conclusions about the impact that the reaction to oil prices has on the distribution of losses across the model space.

Table 9 - Outcome dispersion, different values of the parameters in the process for the real price of oil

	OT rule	S rule	BG rule	H rule	OT rule	S rule	BG rule	H rule
	$\sigma_\xi^2= 80^2$				$\sigma_\xi^2= 90^2$			
g_i	0	0.608	0.073	0.473	0	0.599	0.069	0.451
g_s	0	-0.007	-0.0194	-0.004	0	-0.0071	-0.0195	-0.0041
mean	48.44	38.49	36.73	42.65	55.72	43.15	41.01	48.51
variance	27.18	16.99	14.41	22.16	33.99	20.99	18.05	27.41
median	39.43	32.34	32.52	34.76	44.62	35.20	36.03	38.40
p.w. mean	35.99	32.51	33.94	33.28	40.74	36.71	38.30	37.68
	$\rho = 0.88$				$\rho = 0.95$			
g_i	0	0.609	0.082	0.498	0	0.610	0.050	0.410
g_s	0	-0.0062	-0.0188	-0.0033	0	-0.0094	-0.0204	-0.0061
mean	40.00	33.79	32.27	36.55	96.86	63.63	62.20	75.11
variance	18.94	12.74	9.84	16.60	75.69	47.64	42.72	53.11
median	33.76	29.28	29.30	31.54	72.08	49.10	50.79	53.44
p.w. mean	30.71	28.03	29.66	28.55	65.32	56.31	57.47	58.74

Note: Outcome dispersion analysis for alternative values of the parameters σ_ξ^2 and ρ in (17). For the S, BG and H rules, the new values of the coefficients g_i and g_s in (18), computed using the same procedure as in the baseline scenario, are also reported.

I examined the sensitivity of the outcome dispersion analysis to variations in the magnitude of the volatilities σ_o^2 and $\sigma_{\varepsilon_\delta}^2$ and of the autoregressive coefficient ρ. In each case, I first computed the new S, BG and H rules using the same procedure described in the previous section. Then, I studied the distribution of expected losses generated by these new alternative simple rules, and I compared it with the performance of the original Taylor rule. The results of this exercise are summarized in table 9.

In the baseline case, I set $\sigma_o^2 = 4^2(220)$ and $\sigma_{\varepsilon_\delta}^2 = 4^2(1.9)$, so that $\sigma_\xi^2 = 4^2(441.9) = 84.09^2$. I considered changes in the volatility of the innovations in the process for s_t in the range $\left[\sigma_\xi^2 = 80^2, \sigma_\xi^2 = 90^2\right]$.[18] Table 9 shows that these changes have almost no impact on the coefficients of the optimal simple rules, and that the differences in the distribution of losses that these rules imply remain similar to those obtained in the baseline scenario.

As expected, variations in the autoregressive coefficient ρ have a larger impact on the computed simple rules, in particular on the optimal response of the nominal interest rate to changes in the real price of oil. In the baseline scenario, $\rho = 0.91$; table 9 reports the results for $\rho = 0.88$ and $\rho = 0.95$. In terms of outcome dispersion, the differences between the distribution of losses generated by the OT rule and those originated by the alternative simple rules become larger when changes in the real price of oil are more persistent. Overall, all cases considered in table 9 confirm the result obtained in the baseline scenario that expected losses exhibit the highest mean, variance, median and posterior weighted mean when the OT rule is implemented.

5.2 Action dispersion

In addition to outcome dispersion, policymakers might be interested in investigating the extent to which the reaction to oil prices recommended by the different specifications in the model space is homogeneous. All the policies reported in table 5 suggest an expansionary response to oil price changes, even if the magnitude of this response is different. These rules were computed using the model specification with the highest posterior probability in each class of models, as explained in the previous section. Here, I study the distribution of the optimal policy reaction to a change in the real price of oil across all the specifications included in the model space. As before, optimal policies were obtained by grid search of the parameters in (18) that minimize (22), under the restriction that the long run effect of output and inflation on the nominal interest rate is the same as in the OT rule.

Table 10 provides information on the interest rate response to a change in the real price of oil recommended by the different specifications in the model space. The table focuses on the coefficient g_s in (18) and on the long run effect of oil price changes on the nominal interest

[18] This range corresponds to reasonable values of σ_o^2 and $\sigma_{\varepsilon_\delta}^2$ in the process for the real price of oil, according to the results reported in Rondina (2010).

rate, defined as: $\widetilde{g}_s = g_s/(1 - g_i)$. The results are reported for each class of models and for the entire model space. From this table, it is clear that the recommended response to a change in the real price of oil varies considerably across the model space. In average, the elements of the Blanchard-Gali class suggest a stronger response relative to the model specifications belonging to the other two classes, as evident from rows (1) and (8). The standard deviation of g_s, which measures the contemporaneous reaction of i_t to oil price variations, is higher in the Blanchard-Gali class, while the dispersion of $\widetilde{g}_s$, that is the long run effect of a change in oil prices on i_t, is higher in the Solow and Hamilton classes. This difference is due to the term for interest rate inertia, g_i, which is larger in average in the model specifications belonging to these last two classes. Finally, a number of models, particularly in the Hamilton class, recommend a positive response to changes in the real price of oil. In these specifications, the (positive) impact of oil prices on core CPI inflation is larger than the (negative) impact on the output gap, so that a contractionary rather than expansionary policy is required to contrast the economic consequences of a change in the real price of oil.

Table 10 - Distribution of the optimal response to real oil price changes

	All models		M^S		M^{BG}		M^H	
	g_s	$\widetilde{g}_s$	g_s	$\widetilde{g}_s$	g_s	$\widetilde{g}_s$	g_s	$\widetilde{g}_s$
(1) Mean	-0.0140	-0.001	-0.0080	0.0036	-0.0231	-0.0262	-0.0060	0.0335
(2) Standard deviation	0.0138	0.1109	0.0066	0.1617	0.0149	0.0171	0.0076	0.1253
(3) Minimum	-0.0475	-0.1333	-0.0214	-0.1120	-0.0475	-0.0522	-0.0246	-0.1333
(4) Q1	-0.0208	-0.0333	-0.0126	-0.0313	-0.0359	-0.0414	-0.0123	-0.0209
(5) Median	-0.0135	-0.0210	-0.0089	-0.0195	-0.0210	-0.0242	-0.0041	-0.0084
(6) Q3	-0.0031	-0.0074	-0.0008	-0.0059	-0.0180	-0.0208	0.0001	0.1000
(7) Maximum	0.0203	1.1333	0.0102	1.1333	0.0203	0.0244	0.0056	0.6000
(8) Post. weighted average	-0.0075	-0.0205	-0.0074	-0.0215	-0.0222	-0.0245	-0.0047	-0.0117
(9) N. of models	198		55		87		56	

Notes: 1. Distribution of the recommended values of g_s and $\widetilde{g}_s = g_s/(1 - g_i)$ across all specifications included in the model space.

2. The composition of each class of models is described in Appendix 1.

5.3 Minimax and minimax regret

As last step of the analysis, I examined the policy recommendations of the minimax and minimax regret criteria, defined by (6) and (7) in section 2. This exercise was performed using the OT rule and the S, BG and H rules described in table 5. As previously discussed, the non-Bayesian minimax and minimax regret approaches do not take into account the models' posterior probabilities. Therefore, in this portion of the policy evaluation an equal weight is attached to all the specifications included in the model space.

Table 11 - Minimax analysis

	(1) All models	(2) M^S	(3) M^{BG}	(4) M^H
N. of models	198	55	87	56
Max Loss				
Taylor rule	157.41	78.55	157.41	67.20
S rule	127.63	75.67	127.63	64.32
BG rule	130.24	75.53	130.24	63.72
H rule	132.26	77.02	132.26	65.59
Minimax	S rule	BG rule	S rule	BG rule

Notes: 1. Robust policy rule recommended by the minimax criterion for each class of models and for the entire model space. The minimax criterion is defined by (6).
2. The OT rule is defined by (21) and the S, BG and H rules are defined by (18), with coefficient values as reported in table 5.
3. The composition of each class of models is described in Appendix 1.

Table 12 - Minimax regret analysis

	(1) All models	(2) M^S	(3) M^{BG}	(4) M^H
N. of models	198	55	87	56
Max Regret				
Taylor rule	83.22	8.25	83.22	9.53
S rule	57.03	0.95	57.03	3.79
BG rule	59.65	5.52	59.65	10.25
H rule	57.71	4.00	57.71	5.70
Minimax Regret	S rule	S rule	S rule	S rule

Notes: 1. Robust policy rule recommended by the minimax regret criterion for each class of models and for the entire model space. The minimax regret criterion is defined by (7).
2. The OT rule is defined by (21) and the S, BG and H rules are defined by (18), with coefficient values as reported in table 5.
3. The composition of each class of models is described in Appendix 1.

Table 11 reports the result of the minimax analysis for each class of models and for the entire model space. Across the 198 specifications composing the model space, the policy rule that minimizes the maximum possible loss is the S rule. This is also the case if we only consider the Blanchard-Gali class of models, while in the other two classes the BG rule delivers a (slightly) lower maximum loss. In all sets of models, the OT rule induces the highest maximum loss.

Table 12 reports the policy recommendations of the minimax regret criterion. For each model, regret is defined as the difference between the loss suffered by a policy and the loss under the optimal policy for that specific model. Thus, relative to the minimax criterion, this measure is able to reduce the dominance of those specifications that entail relatively high losses regardless of the selected policy. Table 12 shows that the policy minimizing the maximum regret, in each class of models and in the entire model space, is the S rule. Again, in all sets of models, the OT rule delivers the highest maximum regret.

For the space of model specifications considered in this work, tables 11 and 12 show that the minimax and the minimax regret criteria both recommend the same policy, that is the S rule. This policy is also the one that generates the lowest posterior weighted average loss across the model space, as reported in table 7. Thus, among the policy rules considered in the policy evaluation exercise, the Bayesian model averaging approach and the non-Bayesian minimax and minimax regret criteria agree on the choice of the robust policy under model uncertainty. Moreover, under all measures the least recommended policy is always the original Taylor rule.

5.4 An alternative model space

In the baseline scenario, the definition of the model space was centred on the three different theories that policymakers believe as possibly generating the data. As a consequence, the specifications included in the restricted model space used for the outcome dispersion, action dispersion, minimax and minimax regret analysis were those with the highest posterior probabilities within each class of models. In this section, I investigate whether the results of the policy evaluation exercise would be different under an alternative definition of the model space that puts less emphasis on the theory from which each model specification originates.

The model space considered in this section was defined using the following procedure. Starting from the initial set of $30,720$ specifications, I attached the same initial weight to all of them by assuming a uniform prior of $1/30720$. Then, I selected all the models with posterior probability of at least $1/200 = 0.5\%$ of the model with the highest posterior in the entire model space.[19] In this way, only specifications with high posterior probability in absolute (and not in relative) terms were included in the restricted model space used for the policy evaluation exercise. This procedure selected a total of 96 models, covering 89.27% of the posterior probability.

[19] I decreased the threshold relative to the baseline scenario to include an overall posterior probability comparable with those reported in table 3 for the different classes of models. In any case, the same exercised performed with the threshold of 1% delivers very similar results.

Of these, 89 were part of the original Solow class of models, 1 of the Blanchard-Gali class, and 6 of the Hamilton class. The higher prior attached to the Solow specifications relative to the baseline case is reflected in the composition of the restricted model space, which is almost entirely constituted of models belonging to this class. The specification with the highest posterior probability in this alternative definition of the model space corresponds to the specification with the highest posterior in the Solow class, so its estimated coefficients were already reported in table 4.

Table 13 - Distribution of model losses in the alternative model space

	OT rule	S rule	BG rule	H rule
(1) Mean	42.406	39.065	40.508	39.776
(2) Standard deviation	12.516	11.754	11.0945	12.395
(3) Minimum	21.352	20.344	22.549	19.459
(4) Q1	34.006	30.279	33.281	30.535
(5) Median	40.205	36.521	38.340	37.291
(6) Q3	49.283	45.397	45.701	47.012
(7) Maximum	79.226	75.669	75.527	77.020
(8) Posterior weighted average	38.729	35.379	37.033	36.101
(9) N. of models	96	96	96	96

Notes: 1. Distribution of model losses under different policy rules. The OT rule is defined by (21) and the S, BG and H rules are defined by (18), with coefficient values as reported in table 5.
2. Rows (1) - (7) report basic statistics of the distribution of losses under each policy rule. Row (8) reports the posterior weighted average loss, computed using (23).
3. The models space is composed of 96 models, 89 from the Solow class, 1 from the Blanchard-Gali class, and 6 from the Hamilton class. These models were selected using the procedure described in the main text.

Table 13 provides some summary statistics of the distribution of losses across the new model space for the policy rules that were studied in the original analysis. The S rule corresponds to the policy recommended by the specification with the highest posterior in this new model space. In addition, the Blanchard-Gali and Solow specifications described in table 4, which were used to compute the BG and H rules, are still part of the model space, even in this alternative definition. For this reason, as well as for comparison purposes, the policy evaluation exercise was performed using the same set of policies considered in the baseline model scenario.

The performance of the OT rule in terms of the first two moments of the distribution of losses across the model space is considerably improved in this case. This result was somehow expected, since this policy rule originates high losses particularly in the Blanchard-Gali class of models, which is greatly underrepresented here compared to the baseline scenario (1 specification instead of 87). In addition, the marginal presence of Blanchard-Gali specifications for which, as shown in figure 3, losses exhibit a general tendency to be more volatile, induces a reduction in the standard deviation of losses under all policy rules. Nonetheless, the differences in performance in terms of posterior weighted average loss are almost the same as those reported in table 7. This happens because the Solow specifications dominate the model space in terms of posterior probabilities, in this exercise as well as in the baseline case. Therefore, the differences in the expected losses generated by the selected policies in the two scenarios almost disappear when these are weighted using the models' posteriors.

Figure 5 - Model losses for each policy relative to the Taylor rule in the alternative model space

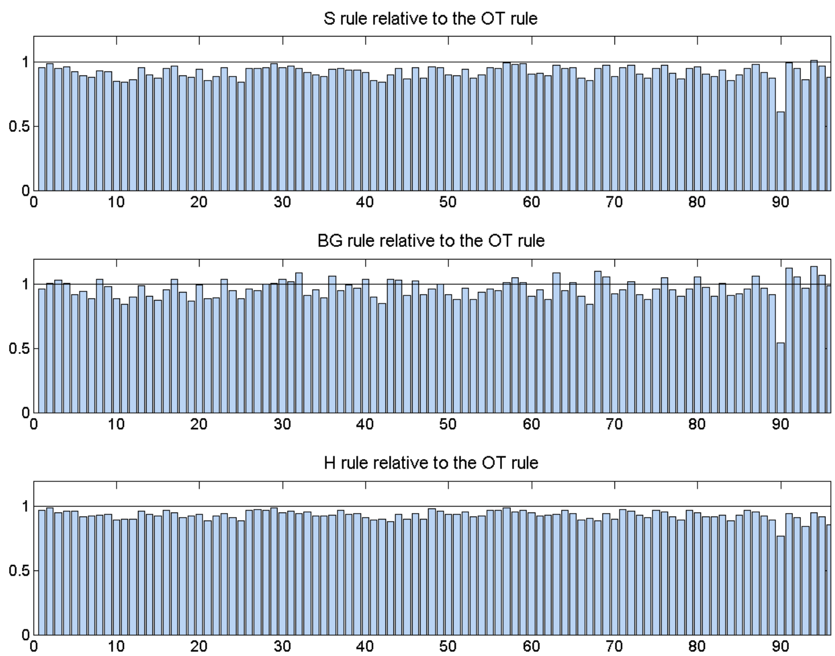

Notes: 1. Each panel reports the ratio between the loss generated by one of the simple policy rules described in table 5 and the loss generated by the original Taylor rule, for each specification in the model space.
2. Model numbers are as follows: specifications from 1 to 89 belong to the Solow class of models, specification 90 belongs to the Blanchard-Gali class, and specifications from 91 to 96 belong to the Hamilton class. These models have been selected using the procedure described in the main text.

Finally, figure 5 provides some additional information about the losses generated by the S, BG and H policy rules relative to the OT rule. This figure shows that, while for some elements of the model space the OT rule is able to outperform the BG rule in terms of model specific losses, this is almost never the case when this rule is compared to the S and H policies. In all, from table 13 and figure 5 we can conclude that even if the OT rule is considerably more comparable to the other policies in this different model space, it still produces the highest average losses, either non-weighted or weighted using the models' posterior probabilities.

In this alternative definition of the model space, based on the posterior weighted average losses reported in table 13 a Bayesian policymaker would select the S rule. On the other hand, the minimax criterion would suggest the BG rule, while the minimax regret approach would recommend again the S rule. Relative to the baseline scenario, only the minimax criterion selects a different policy rule. However, it is clear from line (7) in table 13 that the lead of the BG policy rule over the S rule is minimal, since the maximum loss generated by these two policies is actually almost the same. Therefore, the conclusion that in this environment various approaches point to the S rule as the robust policy choice under model uncertainty can be regarded as valid even under the different definition of model space considered in this section. Furthermore, all the techniques still regard the original Taylor rule as the least robust among the set of policies under study.

6 Concluding remarks

In this paper, I analyzed the problem of a policymaker that is uncertain about the mechanisms through which oil prices affect economic activity. I conducted an empirical study of the likelihood of three alternative theories that have been proposed to explain the effects of oil prices on the economy, and I presented a policy evaluation exercise encompassing a range of techniques that have been developed in the model uncertainty literature. In this environment, I found that according to a number of Bayesian and non-Bayesian measures, the original Taylor rule performs worse than a set of alternative simple rules in which policymakers introduce persistence in the nominal interest rate and respond to changes in the real price of oil. In particular, I showed that allowing the policy rule to react to oil prices is important for controlling the mean and volatility of expected losses across the different specifications considered in the analysis. This result was not obvious. Since the different elements of the model space recommend a contrasting optimal response to oil prices (negative in some cases, positive in others) it could have been possible as well that a policy rule not reacting to changes in the real price of oil performed better than another one imposing a response in one specific direction.

I believe that this work could be extended in a few different directions. First, the policy analysis could be enriched to account for the lack of consensus on the way oil prices should be measured. Indeed, while a part of the literature focused on real oil prices, in levels or differences (see, for instance, Blanchard and Gali, 2007 and Herrera and Pesavento, 2009), other contributions introduced alternative measures of nominal oil price changes (see, among the others, Hamilton, 2003 and Cavallo and Wu, 2009). This issue could be incorporated in the framework proposed in this paper by simply considering the uncertainty on the way oil prices should be defined as an additional form of uncertainty characterizing the model space.

A second extension could be the inclusion of 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 (2005), 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 would decrease demand for some goods and possibly increase demand for others, and it would create incentives for households to postpone their investment activity. However, an oil price decrease would 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 worthy 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. As a first step, the assumption of backward looking expectations could be replaced by the use of 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 used in the main text are the following:

- y_t is the output gap, computed as the difference between real GDP and the CBO estimate of potential GDP, 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 annualized change in the real price of oil. The real price of oil is defined 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 average Federal Funds rate.

The data is quarterly and includes observations from $1973 : I$ to $2008 : II$, with data from $1971 : I$ to $1972 : IV$ used to provide lags. All the data was obtained from the Federal Reserve Bank of St. Louis web site.

The computation of the expected losses defined by (22) requires that policymakers know the value of the parameters in the process for the real price of oil. As explained in the main text, Rondina (2010) estimates these parameters using a MCMC algorithm. Given the results of this related work, I set $\rho = 0.91$, $\sigma_o^2 = 4^2\,(220)$ and $\sigma_\varepsilon^2 = 4^2\,(1.9)$. This implies that $\sigma_\xi^2 = 4^2\,(441.9) = 84.09^2$.

Model labeling

The full model space includes $30,720$ models, $20,480$ for M^S and $5,120$ for M^{BG} and M^H. The numbering of the models is organized as follows:

- models from 1 to $20,480$ are the S class of models;
- models from $20,481$ to $25,600$ are the BG class of models;
- models from $25,601$ to $30,720$ are the H class of models.

The elements of each class of models differ in terms of the variables and the number of lags of each variable included in the output and inflation equations, as specified in table 1. In each class of models, the order in which the lags change is the following:

1. lags of s_t in the inflation equation;
2. lags of i_t in the inflation equation (for the S and H models);
3. lags of π_t in the inflation equation;
4. lags of y_t in the inflation equation;
5. lags of s_t in the output equation;
6. lags of the real interest rate $(i_{t-1} - E_{t-1}(\pi_t))$ in the output equation (for the BG models);
7. lags of unanticipated inflation $(\pi_t - E_{t-1}(\pi_t))$ in the output equation (for the S models);
8. lags of y in the output equation.

In the policy evaluation exercise, I only consider a subset of the initial model space, composed of 55 models for the S class, 87 models for the BG class and 56 models for the H class for a total of 198 model specifications. The process used to select these models was explained in the main text. The numbering of the elements in this restricted model space is as follows:

- models from 1 to 55 are the S class of models;
- models from 56 to 142 are the BG class of models;
- models from 143 to 198 are the H class of models.

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

Table 14 - Model specifications in the restricted Solow class

		Lags output equation			Lags inflation equation			
Number	Original model number	y	π-E(π)	s	y	π	i	s
1	6446	2	2	1	1	3	2	0
2	6456	2	2	1	1	3	4	0
3	6476	2	2	1	1	4	4	0
4	6477	2	2	1	1	4	4	1
5	6536	2	2	1	2	3	4	0
6	6556	2	2	1	2	4	4	0
7	6796	2	2	2	1	4	4	0
8	7726	2	3	1	1	3	2	0
9	7736	2	3	1	1	3	4	0
10	7756	2	3	1	1	4	4	0
11	7836	2	3	1	2	4	4	0
12	9036	2	4	1	1	4	4	0
13	11566	3	2	1	1	3	2	0
14	11567	3	2	1	1	3	2	1
15	11571	3	2	1	1	3	3	0
16	11576	3	2	1	1	3	4	0
17	11577	3	2	1	1	3	4	1
18	11578	3	2	1	1	3	4	2
19	11586	3	2	1	1	4	2	0
20	11596	3	2	1	1	4	4	0
21	11597	3	2	1	1	4	4	1
22	11598	3	2	1	1	4	4	2
23	11646	3	2	1	2	3	2	0
24	11656	3	2	1	2	3	4	0
25	11676	3	2	1	2	4	4	0
26	11677	3	2	1	2	4	4	1
27	11678	3	2	1	2	4	4	2
28	11736	3	2	1	3	3	4	0
29	11756	3	2	1	3	4	4	0
30	11836	3	2	1	4	4	4	0
31	11886	3	2	2	1	3	2	0
32	11896	3	2	2	1	3	4	0
33	11916	3	2	2	1	4	4	0
34	11917	3	2	2	1	4	4	1
35	11996	3	2	2	2	4	4	0
36	12846	3	3	1	1	3	2	0
37	12856	3	3	1	1	3	4	0
38	12857	3	3	1	1	3	4	1
39	12866	3	3	1	1	4	2	0
40	12876	3	3	1	1	4	4	0
41	12877	3	3	1	1	4	4	1
42	12878	3	3	1	1	4	4	2
43	12936	3	3	1	2	3	4	0
44	12956	3	3	1	2	4	4	0
45	12957	3	3	1	2	4	4	1
46	13196	3	3	2	1	4	4	0
47	14156	3	4	1	1	4	4	0
48	16686	4	2	1	1	3	2	0
49	16696	4	2	1	1	3	4	0
50	16716	4	2	1	1	4	4	0
51	16717	4	2	1	1	4	4	1
52	16776	4	2	1	2	3	4	0
53	16796	4	2	1	2	4	4	0
54	17036	4	2	2	1	4	4	0
55	17996	4	3	1	1	4	4	0

Table 15 - Model specifications in the restricted Blanchard-Gali class

		Lags output equation			Lags inflation equation		
Number	Original Model number	y	i- E(π)	s	y	π	s
56	21771	2	1	1	1	3	0
57	21791	2	1	1	2	3	0
58	22091	2	2	1	1	3	0
59	22092	2	2	1	1	3	1
60	22093	2	2	1	1	3	2
61	22096	2	2	1	1	4	0
62	22111	2	2	1	2	3	0
63	22112	2	2	1	2	3	1
64	22113	2	2	1	2	3	2
65	22116	2	2	1	2	4	0
66	22131	2	2	1	3	3	0
67	22151	2	2	1	4	3	0
68	22171	2	2	2	1	3	0
69	22176	2	2	2	1	4	0
70	22191	2	2	2	2	3	0
71	22192	2	2	2	2	3	1
72	22193	2	2	2	2	3	2
73	22196	2	2	2	2	4	0
74	22211	2	2	2	3	3	0
75	22231	2	2	2	4	3	0
76	22411	2	3	1	1	3	0
77	22431	2	3	1	2	3	0
78	22471	2	3	1	4	3	0
79	22491	2	3	2	1	3	0
80	22511	2	3	2	2	3	0
81	23051	3	1	1	1	3	0
82	23071	3	1	1	2	3	0
83	23371	3	2	1	1	3	0
84	23372	3	2	1	1	3	1
85	23373	3	2	1	1	3	2
86	23376	3	2	1	1	4	0
87	23391	3	2	1	2	3	0
88	23392	3	2	1	2	3	1
89	23393	3	2	1	2	3	2
90	23394	3	2	1	2	3	3
91	23396	3	2	1	2	4	0
92	23398	3	2	1	2	4	2
93	23411	3	2	1	3	3	0
94	23413	3	2	1	3	3	2
95	23416	3	2	1	3	4	0
96	23431	3	2	1	4	3	0
97	23432	3	2	1	4	3	1
98	23433	3	2	1	4	3	2
99	23436	3	2	1	4	4	0
100	23451	3	2	2	1	3	0
101	23452	3	2	2	1	3	1
102	23456	3	2	2	1	4	0
103	23471	3	2	2	2	3	0
104	23472	3	2	2	2	3	1
105	23473	3	2	2	2	3	2
106	23476	3	2	2	2	4	0
107	23491	3	2	2	3	3	0
108	23511	3	2	2	4	3	0
109	23531	3	2	3	1	3	0
110	23551	3	2	3	2	3	0
111	23691	3	3	1	1	3	0
112	23692	3	3	1	1	3	1
113	23696	3	3	1	1	4	0
114	23711	3	3	1	2	3	0
115	23712	3	3	1	2	3	1
116	23713	3	3	1	2	3	2
117	23716	3	3	1	2	4	0
118	23731	3	3	1	3	3	0
119	23751	3	3	1	4	3	0
120	23771	3	3	2	1	3	0
121	23791	3	3	2	2	3	0
122	23831	3	3	2	4	3	0
123	24011	3	4	1	1	3	0
124	24031	3	4	1	2	3	0
125	24071	3	4	1	4	3	0
126	24111	3	4	2	2	3	0
127	24651	4	2	1	1	3	0
128	24652	4	2	1	1	3	1
129	24656	4	2	1	1	4	0
130	24671	4	2	1	2	3	0
131	24672	4	2	1	2	3	1
132	24673	4	2	1	2	3	2
133	24676	4	2	1	2	4	0
134	24691	4	2	1	3	3	0
135	24711	4	2	1	4	3	0
136	24731	4	2	2	1	3	0
137	24751	4	2	2	2	3	0
138	24971	4	3	1	1	3	0
139	24991	4	3	1	2	3	0
140	25071	4	3	2	2	3	0
141	25291	4	4	1	1	3	0
142	25311	4	4	1	2	3	0

Table 16 - Model specifications in the restricted Hamilton class

		Lags output equation		Lags inflation equation						Lags output equation		Lags inflation equation			
Number	Original model number	y	s	y	π	i	s	Number	Original model number	y	s	y	π	i	s
143	26926	2	1	1	3	2	0	171	28217	3	1	1	3	4	1
144	26927	2	1	1	3	2	1	172	28218	3	1	1	3	4	2
145	26931	2	1	1	3	3	0	173	28226	3	1	1	4	2	0
146	26936	2	1	1	3	4	0	174	28236	3	1	1	4	4	0
147	26937	2	1	1	3	4	1	175	28237	3	1	1	4	4	1
148	26938	2	1	1	3	4	2	176	28238	3	1	1	4	4	2
149	26946	2	1	1	4	2	0	177	28286	3	1	2	3	2	0
150	26956	2	1	1	4	4	0	178	28296	3	1	2	3	4	0
151	26957	2	1	1	4	4	1	179	28316	3	1	2	4	4	0
152	26958	2	1	1	4	4	2	180	28317	3	1	2	4	4	1
153	27006	2	1	2	3	2	0	181	28318	3	1	2	4	4	2
154	27016	2	1	2	3	4	0	182	28376	3	1	3	3	4	0
155	27036	2	1	2	4	4	0	183	28396	3	1	3	4	4	0
156	27037	2	1	2	4	4	1	184	28476	3	1	4	4	4	0
157	27038	2	1	2	4	4	2	185	28526	3	2	1	3	2	0
158	27096	2	1	3	3	4	0	186	28536	3	2	1	3	4	0
159	27116	2	1	3	4	4	0	187	28556	3	2	1	4	4	0
160	27246	2	2	1	3	2	0	188	28557	3	2	1	4	4	1
161	27256	2	2	1	3	4	0	189	28616	3	2	2	3	4	0
162	27276	2	2	1	4	4	0	190	28636	3	2	2	4	4	0
163	27277	2	2	1	4	4	1	191	28876	3	3	1	4	4	0
164	27336	2	2	2	3	4	0	192	29486	4	1	1	3	2	0
165	27356	2	2	2	4	4	0	193	29496	4	1	1	3	4	0
166	27596	3	3	1	4	4	0	194	29516	4	1	1	4	4	0
167	28206	3	1	1	3	2	0	195	29517	4	1	1	4	4	1
168	28207	3	1	1	3	2	1	196	29576	4	1	2	3	4	0
169	28211	3	1	1	3	3	0	197	29596	4	1	2	4	4	0
170	28216	2	1	1	3	4	0	198	29836	4	2	1	4	4	0

Appendix 2

Analysis of the policy responses implied by the policy rules described in table 5

This Appendix provides a more in depth analysis of the policy responses implied by the simple rules reported in table 5. Figure 6 investigates the response of output, Core CPI inflation and the Federal Funds rate to a 10% increase in the real price of oil in each of the models described in table 4, when the policymaker implements the policy rule suggested by that specific model. In each panel, 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 (21).

Figure 6 - Impulse responses: simple rules and Taylor rule

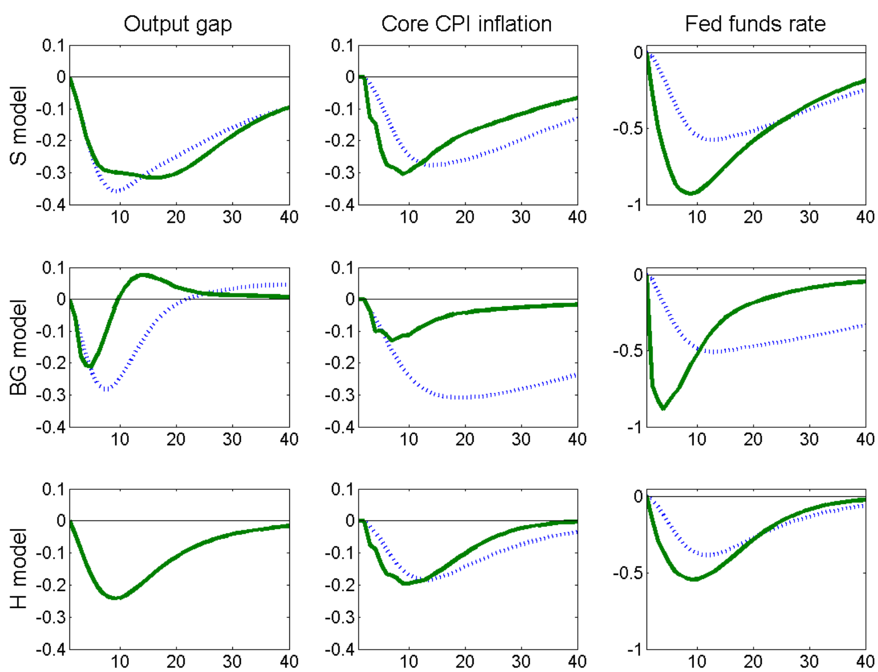

Note: Response of the output gap, core CPI inflation and the Federal funds rate to a 10% increase in the real price of oil. The first column reports the output gap, the second column core CPI inflation and the last column the Federal funds rate. Each row represents a different model and relative policy rule. In each panel, the response of the variable of interest under the selected simple rule (continuous line) is compared to the response under the OT rule described by (21) (dashed line).

All policies recommend a more expansionary response to the change in the real price of oil than the OT rule. In the Hamilton model, which is the one implying the mildest reaction to oil prices, the pattern of all variables is actually quite similar to the impulse-responses that would be generated by implementing the Taylor rule. In the Solow model, the differences between the S and the OT rule are more pronounced, particularly in the behavior of core CPI inflation that returns faster to its the pre-shock level when the S rule is implemented. Lastly, in the Blanchard-Gali model, the BG rule is able to contrast the effects of the oil price increase on the output gap, and is generally a lot more effective than the OT rule in bringing the variables of interest back to their original values.

The next three figures explore the response of output, Core CPI inflation and the Federal Funds rate to a 10% increase in the real price of oil in the models described in table 4, when the policymaker implements one of the policy rules described in table 5.

Figure 7 - Impulse responses: S rule and Taylor rule

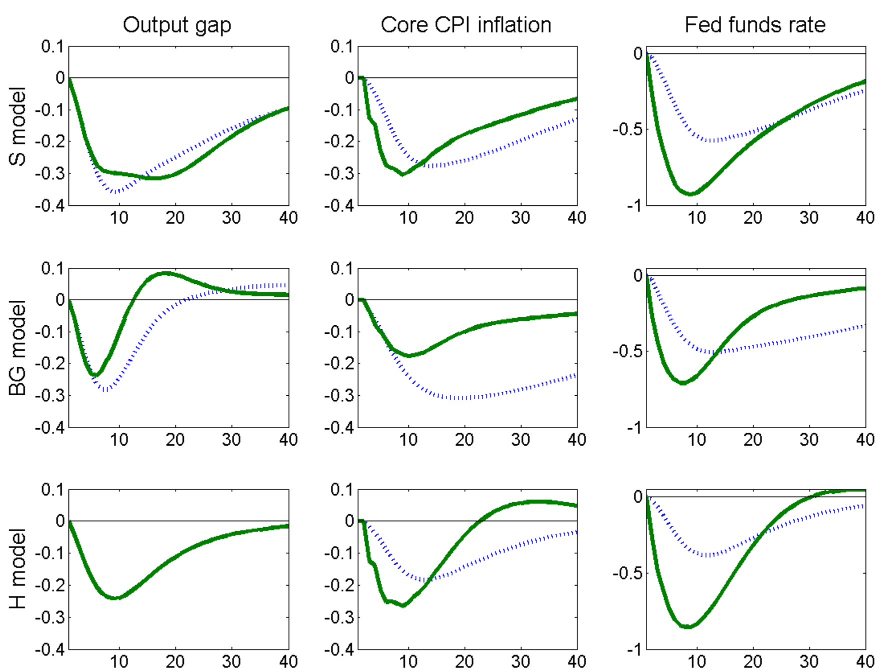

Note: Response of the output gap, core CPI inflation and the Federal funds rate to a 10% increase in the real price of oil. The first columns reports the output gap, the second column core CPI inflation and the last column the Federal funds rate. Each row represents a different model. In each panel, the response of the variable of interest under the S policy rule (continuous line) is compared to the response under the OT rule described by (21) (dashed line).

Figure 7 compares the patterns of the variables of interest generated by the S rule to their behavior under the OT rule. For the Solow model, this exercise is the same as the one performed in figure 6. In the BG model, the S rule is still more effective than the OT rule in contrasting the effects of the increase in oil prices, but the variables return to their original values more slowly compared to the case in which the BG rule is implemented. In the Hamilton model, the S rule generates a stronger response of the nominal interest rate relative to both the OT

rule and the H rule (depicted in figure 6), which leads to more ample variations in core CPI inflation.

Figure 8 - Impulse responses: BG rule and Taylor rule

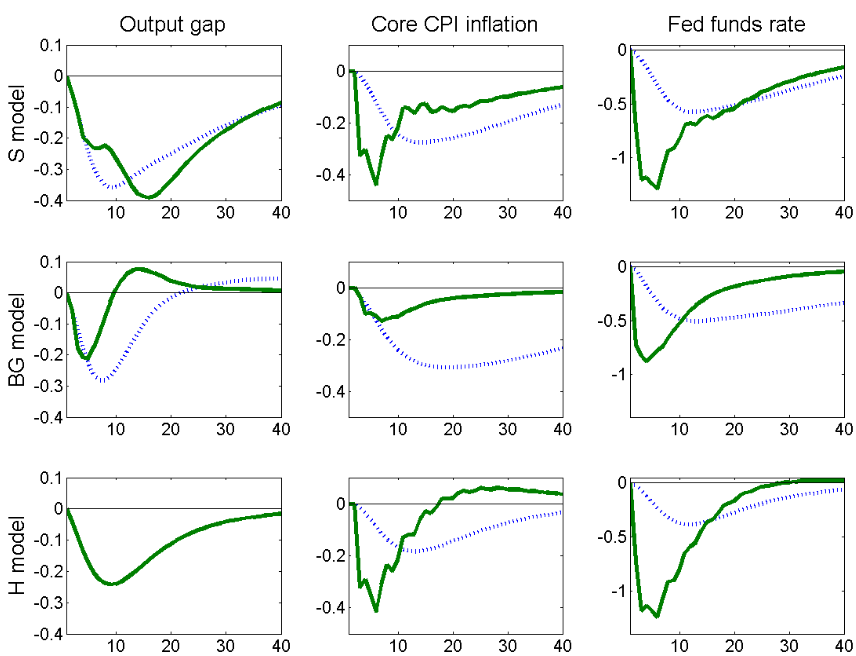

Note: Response of the output gap, core CPI inflation and the Federal funds rate to a 10% increase in the real price of oil. The first columns reports the output gap, the second column core CPI inflation and the last column the Federal funds rate. Each row represents a different model. In each panel, the response of the variable of interest under the BG policy rule (continuous line) is compared to the response under the OT rule described by (21) (dashed line).

Figure 8 studies the patterns of the variables of interest under the BG rule. In the Solow and Hamilton models, the stronger reaction to a change in the real price of oil recommended by the BG rule implies a much larger decrease of the Federal funds rate relative to the OT rule, and to the S and H rules as well (see figures 6 and 7). The interactions between this decrease and the reduction in output gap caused by the increase in oil prices make the pattern of core CPI inflation very volatile and hectic in both models. Nonetheless, in all models the length of the recovery from the oil price increase does not change much relative to the previous exercises.

Finally, figure 9 reports the responses of the output gap, core CPI inflation and Federal funds rate when the H rule is implemented. Given that the reaction to a change in the real price of oil recommended by the H rule is quite small, it is not surprising that, in all models, the patterns of the variables of interest are very similar to those generated by the OT rule.

Figure 9 - Impulse response: H rule and Taylor rule

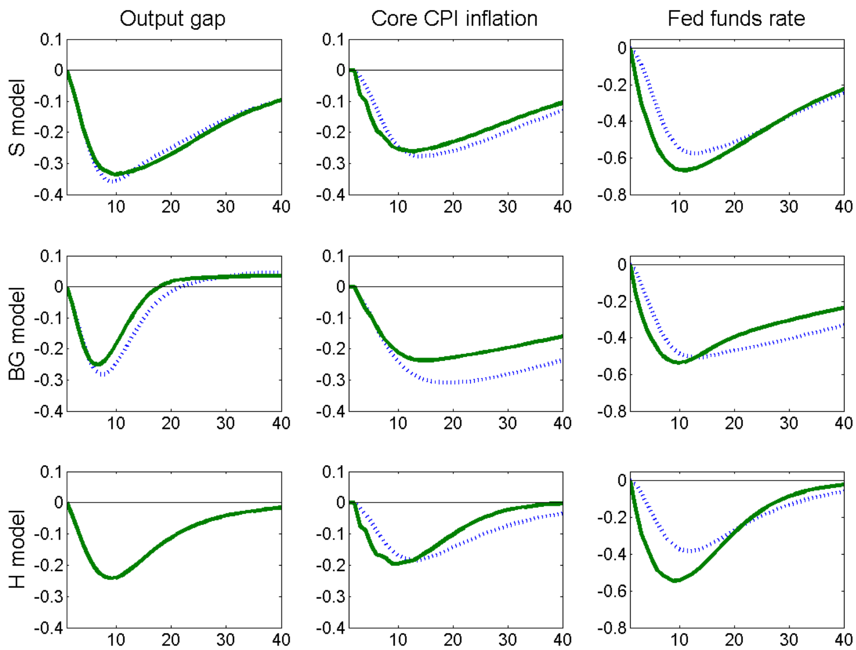

Note: Response of the output gap, core CPI inflation and the Federal funds rate to a 10% increase in the real price of oil. The first columns reports the output gap, the second column core CPI inflation and the last column the federal funds rate. Each row represents a different model. In each panel, the response of the variable of interest under the H policy rule (continuous line) is compared to the response under the OT rule described by (21) (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] Cavallo, M., Wu, T., 2009. Measuring oil-price shocks using market-based information. Working Papers 0905, Federal Reserve Bank of Dallas.

[10] Clark, T. E., Terry, S. J., 2010. Time Variation in the Inflation Passthrough of Energy Prices. Journal of Money, Credit and Banking, vol. 42, no. 7, October, pp. 1419–1433.

[11] Cogley, T., De Paoli, B., Matthes, C., Nikolov, K., Yates, T., 2009. A Bayesian Approach to Optimal Monetary Policy with Parameter and Model Uncertainty. Mimeo, department of economics, New York University.

[12] 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.

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

[14] 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.

[15] 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.

[16] 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.

[17] Hamilton, J.D., 1983. Oil and the Macroeconomy since World War II. Journal of Political Economy, vol. 91, pp. 228-248.

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

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

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

[21] 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.

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

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

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

[25] Herrera, A., Pesavento, E., 2009. Oil Price Shocks, Systematic Monetary Policy, And The "Great Moderation". Macroeconomic Dynamics, vol. 13, n. 1, February, pp. 107-137.

[26] Kilian, L., 2008. The Economic Effects of Energy Price Shocks. Journal of Economic Literature, vol. 46, n.4, December, pp. 871-909.

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

[28] King, R. G., Stock, J. H., Watson, M. W., 1995. Temporal Instability of the Unemployment-Inflation Relationship. Federal Reserve Bank of Chicago Economic Perspectives, XIX, 2–12.

[29] 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.

[30] Levin, A., Wieland, V., Williams, J., 1998. Robustness of simple monetary policy rules under model uncertainty. National Bureau of Economic Research Working Paper No. 6570.

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

[32] Madigan, D., Raftery, A.E., 1994. Model selection and accounting for model uncertainty in graphical models using Occam's window. Journal of the American Statistical Association 89, 1535–1546.

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

[34] 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.

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

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

[37] Pindyck, R.S., 1999. The Long-Run Evolutions of Energy Prices. The Energy Journal, vol. 20, n.2, pp. 1-28.

[38] Primiceri, G.E., 2006. Why inflation rose and fell: Policymakers' beliefs and US postwar stabilization policy. The Quarterly Journal of Economics, vol. 121, August, pp. 867-901.

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

[40] Rondina, F., 2010. The role of model uncertainty and learning in the U.S. postwar policy response to oil prices. UFAE and IAE Working Papers n. 834.10, June.

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

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

[43] 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.

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

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

[46] 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.

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

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