Journal of Economic Dynamics & Control 35 (2011) 2186–2212

Contents lists available at ScienceDirect

Journal of Economic Dynamics & Control

journal homepage: www.elsevier.com/locate/jedc

A Bayesian approach to optimal monetary policy with parameter and model uncertainty

Timothy Cogley [a], Bianca De Paoli [b], Christian Matthes [c], Kalin Nikolov [d,*], Tony Yates [b]

[a] New York University, United States
[b] Bank of England, United Kingdom
[c] Universitat Pompeu Fabra and Barcelona GSE, Spain
[d] European Central Bank, Germany

ARTICLE INFO

Article history:
Received 30 March 2010
Accepted 27 January 2011
Available online 23 February 2011

JEL classification:
E52
C11
C44
C54

Keywords:
Monetary policy
Bayesian analysis
Statistical decision theory
Quantitative policy modeling

ABSTRACT

This paper undertakes a Bayesian analysis of optimal monetary policy for the U.K. We estimate a suite of monetary-policy models that include both forward- and backward-looking representations as well as large- and small-scale models. We find an optimal simple Taylor-type rule that accounts for both model and parameter uncertainty. For the most part, backward-looking models are highly fault tolerant with respect to policies optimized for forward-looking representations, while forward-looking models have low fault tolerance with respect to policies optimized for backward-looking representations. In addition, backward-looking models often have lower posterior probabilities than forward-looking models. Bayesian policies therefore have characteristics suitable for inflation and output stabilization in forward-looking models.

1. Introduction

Central bankers frequently emphasize the importance of uncertainty in shaping monetary policy (e.g. see Greenspan, 2004; King, 2004). Uncertainty takes many forms. The central bank must act in anticipation of future conditions, which are affected by shocks that are currently unknown. In addition, because economists have not formed a consensus about the best way to model the monetary transmission mechanism, policy makers must also contemplate alternative theories with distinctive operating characteristics. Finally, even economists who agree on a modeling strategy sometimes disagree about the values of key parameters. Central bankers must therefore also confront parameter uncertainty within macroeconomic models.

A natural way to address these issues is to regard monetary policy as a Bayesian decision problem. As noted by Brock et al. (2003), a Bayesian approach is promising because it seamlessly integrates econometrics and decision theory. Thus, we can use Bayesian econometric methods to assess various sources of uncertainty and incorporate the results as an input to a decision problem.

* Correspondence to: Bank of England, Monetary Analysis Department, 33 Threadneedle Street, London, United Kingdom. Tel.:+442076015408.
E-mail addresses: tim.cogley@nyu.edu (T. Cogley), bianca.depaoli@bankofengland.co.uk (B. De Paoli), christian.matthes@upf.edu (C. Matthes), kalin.nikolov@ecb.int (K. Nikolov), tony.yates@bankofengland.co.uk (T. Yates).

0165-1889/$ - see front matter
doi:10.1016/j.jedc.2011.02.006

Our aim in this paper is to consider how monetary policy should be conducted in the face of multiple sources of uncertainty, including model and parameter uncertainty as well as uncertainty about future shocks. We apply Bayesian methods root and branch to a suite of macroeconomic models estimated on U.K. data, and we use the results to devise a simple, optimal monetary-policy rule.

1.1. The method in more detail

Just to be clear, we take two shortcuts relative to a complete Bayesian implementation. First, we neglect experimentation. Under model and/or parameter uncertainty, a Bayesian policy maker has an incentive to vary the policy instrument in order to generate information about unknown parameters and model probabilities. In the context of monetary policy, however, a number of recent studies suggest that experimental motives are weak and that 'adaptive optimal policies' (in the language of Svensson and Williams, 2008a) well approximate fully optimal, experimental policies.[1] Because of that, and also because many central bankers are averse to experimentation, our goal is to formulate an optimal non-experimental rule.

We also restrict attention to a simple rule, i.e. one involving a relatively small number of arguments as opposed to the complete state vector. This is for tractability as well as for transparency. For a Bayesian decision problem with multiple models, the fully optimal decision rule would involve the complete state vector for all the models under consideration. That would complicate our calculations a great deal. Some economists also argue that simple rules constitute more useful communication tools. For example, Woodford (1999) writes that "a simple feedback rule would make it easy to describe the central bank's likely future conduct with considerable precision, and verification by the private sector of whether such a rule is actually being followed should be straightforward as well." Thus, we restrict policy to follow Taylor-like rules.

With those simplifications in mind, our goal is to choose the parameters of a Taylor rule to minimize expected posterior loss. Suppose ϕ represents the policy-rule parameters and that $l_i(\phi,\theta_i)$ represents expected loss conditional on a particular model i and a calibration of its parameters θ_i. Typically $l_i(\phi,\theta_i)$ is a discounted quadratic loss function that evaluates uncertainty about future shocks. One common approach in the literature is to choose ϕ to minimize $l_i(\phi,\theta_i)$. This delivers a simple optimal rule for a particular model and calibration, but it neglects parameter and model uncertainty.

To incorporate parameter uncertainty within model i, we must first assess how much uncertainty there is. This can be done by simulating the model's posterior distribution, $p(\theta_i|Y,M_i)$, where M_i indexes model i, and Y represents current and past data on variables relevant for that model. Methods for Bayesian estimation of DSGE models were pioneered by Schorfheide (2000) and Smets and Wouters (2003) and are reviewed by An and Schorfheide (2007). If model i were the only model under consideration, expected loss would be

$$l_i(\phi) = \int l_i(\phi,\theta_i)p(\theta_i|Y,M_i)\,d\theta_i. \tag{1}$$

This integral might seem daunting, but it can be approximated by averaging across draws from the posterior simulation. Assuming evenly weighted draws from the posterior, expected loss is

$$l_i(\phi) \approx N^{-1}\sum_{j=1}^{N} l_i(\phi,\theta_{ij}), \tag{2}$$

where N represents the number of Monte Carlo draws and θ_{ij} is the jth draw for model i. A policy rule robust to parameter uncertainty within model i can be found by choosing ϕ to minimize $l_i(\phi)$.

This is a step forward, but it still neglects model uncertainty. To incorporate multiple models, we attach probabilities to each and weigh their implications in accordance with those probabilities. Posterior model probabilities depend on prior beliefs and on their fit to the data. Suppose that $p(M_i)$ is the policy-makers prior probability on model i, that $p(\theta_i|M_i)$ summarizes his prior beliefs about the parameters of that model, and that $p(Y|\theta_i,M_i)$ is the model's likelihood function.[2] According to Bayes' theorem, the posterior model probability is

$$p(M_i|Y) \propto p(Y|M_i)p(M_i), \tag{3}$$

where

$$p(Y|M_i) = \int p(Y|\theta_i,M_i)p(\theta_i|M_i)\,d\theta_i \tag{4}$$

is the marginal likelihood or marginal data density. The latter can also be approximated numerically using output of the posterior simulation; see An and Schorfheide for details. To account for model uncertainty, we average $l_i(\phi)$ across models

[1] E.g. see Cogley et al. (2007, 2008), and Svensson and Williams (2007a, 2007b, 2008a, 2008b).

[2] For simplicity, we assume that Y is common across models, but that is unnecessary. A technical appendix posted online at http://homepages.nyu.edu/~tc60 describes the more realistic case in which the list of variables differs across model.