Journal of Monetary Economics 55 (2008) S80–S96

Contents lists available at ScienceDirect

Journal of Monetary Economics

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

Learning, expectations formation, and the pitfalls of optimal control monetary policy ☆

Athanasios Orphanides [a], John C. Williams [b,*]

[a] *Central Bank of Cyprus, 80 Kennedy Avenue, 1076 Nicosia, Cyprus*
[b] *Federal Reserve Bank of San Francisco, 101 Market Street, San Francisco, CA 94105, USA*

A R T I C L E   I N F O

Article history:
Received 29 November 2007
Received in revised form
6 August 2008
Accepted 6 August 2008
Available online 22 August 2008

JEL Classification System:
E52

Keywords:
Rational expectations
Robust control
Model uncertainty

A B S T R A C T

The optimal control approach to monetary policy has garnered increased attention in recent years. Optimal control policies, however, are designed for the specific features of a particular model and therefore may not be robust to model misspecification. One important source of potential misspecification is how agents form expectations. Specifically, whether they know the complete structure of the model as assumed in rational expectations or learn using a forecasting model that they update based on incoming data. Simulations of an estimated model of the U.S. economy show that the optimal control policy derived under the assumption of rational expectations can perform poorly when agents learn. The optimal control approach can be made more robust to learning by deemphasizing the stabilization of real economic activity and interest rates relative to inflation in the central bank loss function. That is, robustness to learning provides an incentive to employ a "conservative" central banker. In contrast to optimal control policies, two types of simple monetary policy rules from the literature that have been found to be robust to model misspecification in other contexts are shown to be robust to learning.

1. Introduction

For nearly as long as macroeconomic models have existed, economists have been drawn to the idea of applying optimal control theory to the problem of monetary policy (see Chow, 1976, for an early example). Support for the use of optimal control for policy has waxed and waned, reflecting in part swings in economists' confidence in macroeconometric models. Recently, interest in the optimal control approach to monetary policy among academics and central banks has re-emerged, as spelled out in contributions by Svensson (2002), Svensson and Woodford (2003), Woodford (2003), Giannoni and Woodford (2005), and others. As described in Svensson and Tetlow (2006), analytical and computational advances now make it possible to operationalize this approach using the Federal Reserve Board's large-scale nonlinear macroeconomic model and other models.

☆ This paper was prepared for the FRB Dallas Conference, "John Taylor's Contributions to Monetary Theory and Policy," October 12–13, 2007. We thank Jim Bullard, Richard Dennis, Bob Hall, Andy Levin, Julio Rotemberg, David Stockman, Volker Wieland, Mike Woodford, Hakan Yilmazkuday, and an anonymous referee for helpful comments and suggestions. The opinions expressed are those of the authors and do not necessarily reflect the views of the Central Bank of Cyprus, the Governing Council of the European Central Bank, or the management of the Federal Reserve Bank of San Francisco.

* Corresponding author. Tel.: +1 415 974 2240.
E-mail addresses: Athanasios.Orphanides@centralbank.gov.cy (A. Orphanides), John.C.Williams@sf.frb.org (J.C. Williams).

0304-3932/$ - see front matter
doi:10.1016/j.jmoneco.2008.08.002

One potential shortcoming of the optimal control approach is that it ignores uncertainty about the specification of the model. Although in principle one could incorporate various types of uncertainty into the analysis of optimal policy, in practice this is infeasible given current methods and computational power. As a result, existing optimal control policy analysis is done using a single reference model. Levin and Williams (2003) found that optimal control policies can perform very poorly if the central bank's reference model is badly misspecified. Given the prominence accorded to optimal control in the monetary policy literature and increasingly at central banks, it seems an especially propitious moment to examine more closely the robustness properties of optimal control and other monetary policies when the reference model may be misspecified.

This paper examines the robustness of optimal control policies derived under the assumption of rational expectations to potential model misspecification with respect to how private agents form expectations.[1] This approach takes seriously the informational problems facing economic agents, which may cause expectations to deviate from those implied by the model of the economy they inhabit (see Taylor, 1975, for an early analysis of this issue and Sargent, 2007, for a recent discussion). Evidence that survey measures of expectations are inefficient and display significant dispersion at each point in time (see, for example, Mankiw et al., 2003; Williams, 2003a; D'Amico and Orphanides, 2006) suggests that the process of expectations formation is an important source of model uncertainty. As an alternative to rational expectations, agents are assumed to learn about the economy by reestimating a forecasting model as new data become available. A range of plausible learning models are considered, each of which yields good forecasts in our model economy and represents a relatively modest departure from the rational expectations benchmark. Unlike the existing literature in this area, including Orphanides and Williams (2007a, b), which has focused exclusively on simple monetary policy rules that are generalizations of the Taylor (1993) rule, this paper analyzes the robustness of optimal control policies to learning.

The optimal control policy derived assuming rational expectations can perform poorly in an estimated model of the U.S. economy when agents do not possess perfect knowledge of the economy but instead must learn. To evaluate the optimal policy, the central bank is assumed to minimize the weighted unconditional variances of the deviation of the inflation rate from its target, the deviation of the unemployment rate from the natural rate, and the first difference of the nominal interest rate. The optimal control policy brings inflation gradually back to target while smoothing swings in the unemployment rate and interest rate. This policy works extremely well when private expectations are perfectly aligned with those implied by rational expectations; however, if agents are learning, expectations can deviate from those implied by rational expectations, and the finely tuned optimal control policy can go awry. In particular, by implicitly assuming that inflation expectations are always well anchored, the optimal control policy responds insufficiently strongly to movements in inflation, which results in highly persistent and large deviations of the inflation rate from its target.

The lack of robustness of the optimal policy derived under rational expectations can be corrected by assigning the central bank a policy objective that places significantly less weight on stabilization of the real economy and interest rates relative to inflation than that implied by society's true objective function. Thus, robustness to learning provides a rationale for hiring a "conservative" central banker (in the parlance of Rogoff, 1985). The optimal bias to the weights in the loss function are quite large, implying that the central bank should act as if it put up to 10 times greater weight on inflation than implied by society's true loss function. The optimal control policy derived under the biased objective function responds more strongly to inflation, anchoring inflation expectations and performing well under the models of learning that we consider.

The performance of the optimal control policy is then compared to two types of simple monetary policy rules that have been found to be robust to model uncertainty in the literature. The first is a forward-looking version of a Taylor-type rule, which Levin et al. (2003) found to perform very well in a number of estimated rational expectations models. The second is a rule proposed by Orphanides and Williams (2007a) that differs from the first rule in that policy responds to the change in the measure of economic activity, rather than to the level. This type of rule has been shown to perform very well in a counterfactual analysis of monetary policy during 1996–2003 undertaken by Tetlow (2006). Both rules are parsimonious in that they have only two free parameters. These two simple monetary policy rules, where the parameters are chosen to be optimal under rational expectations, perform very well under learning. Indeed, these rules outperform the optimal control policies derived under rational expectations with the society's true loss when agents learn, and yield comparable performance to the optimal control policy derived under the biased objective function.

The remainder of the paper is organized as follows. Section 2 describes the model and reports the estimation results. Section 3 describes the central bank objective and optimal control policy. Section 4 describes the models of expectations formation. Section 5 discusses the simulation methods. Section 6 reports and analyzes economic outcomes for optimal control policies under rational expectations and under learning. Section 7 examines the performance of simple monetary policy rules, and Section 8 concludes.

[1] We follow the methodology advocated by McCallum (1988), Taylor (1993) and implemented in numerous papers, including Taylor (1999), Levin et al. (1999, 2003), Orphanides and Williams (2002, 2007a, b), Brock et al. (2007). An alternative approach that is complementary to that taken in this paper is robust control (Hansen and Sargent, 2007). These methods are best suited to situations where deviations from the reference model are thought to be relatively modest.