

NBER WORKING PAPER SERIES

OPTIMAL MONETARY POLICY UNDER UNCERTAINTY IN DSGE MODELS:
A MARKOV JUMP-LINEAR-QUADRATIC APPROACH

Lars E.O. Svensson
Noah Williams

Working Paper 13892
<http://www.nber.org/papers/w13892>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2008

The authors thank James Bullard, Timothy Cogley, and Andrew Levin for comments on an earlier paper of ours which helped inspire this paper, and Carl Walsh for comments on this paper. The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of any other member of the Executive Board of Sveriges Riksbank. Financial support from the Central Bank of Chile and the National Science Foundation is gratefully acknowledged. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2008 by Lars E.O. Svensson and Noah Williams. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Optimal Monetary Policy under Uncertainty in DSGE Models: A Markov Jump-Linear-Quadratic Approach

Lars E.O. Svensson and Noah Williams

NBER Working Paper No. 13892

March 2008

JEL No. E42,E52,E58

ABSTRACT

We study the design of optimal monetary policy under uncertainty in a dynamic stochastic general equilibrium models. We use a Markov jump-linear-quadratic (MJLQ) approach to study policy design, approximating the uncertainty by different discrete modes in a Markov chain, and by taking mode-dependent linear-quadratic approximations of the underlying model. This allows us to apply a powerful methodology with convenient solution algorithms that we have developed. We apply our methods to a benchmark New Keynesian model, analyzing how policy is affected by uncertainty, and how learning and active experimentation affect policy and losses.

Lars E.O. Svensson
Sveriges Riksbank
SE-103 37 Stockholm
Sweden
and NBER
svensson@princeton.edu

Noah Williams
Department of Economics
001 Fisher Hall
Princeton University
Princeton, NJ 08544-1921
and NBER
noahw@princeton.edu

1 Introduction

In previous work, Svensson and Williams [16] and [17], we have developed methods to study optimal policy in Markov jump-linear-quadratic (MJLQ) models with forward-looking variables: models with conditionally linear dynamics and conditionally quadratic preferences, where the matrices in both preferences and dynamics are random. In particular, each model has multiple “modes,” a finite collection of different possible values for the matrices, whose evolution is governed by a finite-state Markov chain. In our previous work, we have discussed how these modes could be structured to capture many different types of uncertainty relevant for policymakers. Here we put those suggestions into practice. First, we briefly discuss how an MJLQ model can be derived as a mode-dependent linear-quadratic approximation of an underlying nonlinear model. Then, we apply our methods to a simple empirical mode-dependent New Keynesian model of the U.S. economy, a variant of a model by Lindé [11].

In a first paper, Svensson and Williams [16], we studied optimal policy design in MJLQ models when policymakers can or cannot observe the current mode, but we abstracted from any learning and inference about the current mode. Although in many cases the optimal policy under no learning (NL) is not a normatively desirable policy, it serves as a useful benchmark for our later policy analyses. In a second paper, Svensson and Williams [17], we focused on learning and inference in the more relevant situation, particularly for the model-uncertainty applications which interest us, in which the modes are not directly observable. Thus, decision makers must filter their observations to make inferences about the current mode. As in most Bayesian learning problems, the optimal policy thus typically includes an experimentation component reflecting the endogeneity of information. This class of problems has a long history in economics, and it is well-known that solutions are difficult to obtain. We developed algorithms to solve numerically for the optimal policy.¹ Due to the curse of dimensionality, the Bayesian optimal policy (BOP) is only feasible in relatively small models. Confronted with these difficulties, we also considered *adaptive* optimal policy (AOP).² In this case, the policymaker in each period does update the probability distribution of the current

¹ In addition to the classic literature (on such problems as a monopolist learning its demand curve), Wieland [19]-[20] and Beck and Wieland [1] have recently examined Bayesian optimal policy and optimal experimentation in a context similar to ours but without forward-looking variables. Tesfaselassie, Schaling, and Eijffinger [18] examine passive and active learning in a simple model with a forward-looking element in the form of a long interest rate in the aggregate-demand equation. Ellison and Valla [8] and Cogley, Colacito, and Sargent [4] study situations like ours but where the expectational component is as in the Lucas-supply curve ($E_{t-1}\pi_t$, for example) rather than our forward-looking case ($E_t\pi_{t+1}$, for example). More closely related to our present paper, Ellison [7] analyzes active and passive learning in a New Keynesian model with uncertainty about the slope of the Phillips curve.

² What we call optimal policy under no learning, adaptive optimal policy, and Bayesian optimal policy has in the literature also been referred to as myopia, passive learning, and active learning, respectively.

mode in a Bayesian way, but the optimal policy is computed each period under the assumption that the policymaker will not learn in the future from observations. In our setting, the AOP is significantly easier to compute, and in many cases provides a good approximation to the BOP. Moreover, the AOP analysis is of some interest in its own right, as it is closely related to specifications of adaptive learning which have been widely studied in macroeconomics (see Evans and Honkapohja [9] for an overview). Further, the AOP specification rules out the experimentation which some may view as objectionable in a policy context.³

In this paper, we apply our methodology to study optimal monetary-policy design under uncertainty in dynamic stochastic general equilibrium (DSGE) models. We begin by summarizing the main findings from our previous work, leading to implementable algorithms for analyzing policy in MJLQ models. We then turn to analyzing optimal policy in DSGE models. To quantify the gains from experimentation we focus on a small empirical benchmark New Keynesian model. In this model we compare and contrast optimal policies under no learning, AOP, and BOP. We analyze whether learning is beneficial—it is not always so, a fact which at least partially reflects our assumption of symmetric information between the policymakers and the public—and then quantify the additional gains from experimentation.⁴

Since we typically find that the gains from experimentation are small, we focus in the rest of the paper on our adaptive optimal policy which shuts down the experimentation channel. As the AOP is much easier to compute, this allows us to work with much larger and more empirically relevant policy models. In the latter part of the paper, we analyze one such model, an estimated forward-looking model which is a mode-dependent variant of Lindé [11]. There, we focus on how optimal policy should respond to uncertainty about the degree to which agents are forward-looking, and we show that there are substantial gains from learning in this framework.

³ In addition, AOP is useful for technical reasons as it gives us a good starting point for our more intensive numerical calculations in the BOP case.

⁴ In addition to our own previous work, MJLQ models have been widely studied in the control-theory literature for the special case when the model modes are observable and there are no forward-looking variables (see Costa, Fragoso, and Marques [5] (henceforth CFM) and the references therein). (do Val and Başar [6] provide an application of an adaptive-control MJLQ problem in economics.) More recently, Zampolli [22] has used such an MJLQ model to examine monetary policy under shifts between regimes with and without an asset-market bubble. Blake and Zampolli [3] provide an extension of the MJLQ model with observable modes to include forward-looking variables and present an algorithm for the solution of an equilibrium resulting from optimization under discretion. Svensson and Williams [16] provide a more general extension of the MJLQ framework with forward-looking variables and present algorithms for the solution of an equilibrium resulting from optimization under commitment in a timeless perspective as well as arbitrary time-varying or time-invariant policy rules, using the recursive saddlepoint method of Marcet and Marimon [12]. They also provide two concrete examples: an estimated backward-looking model (a three-mode variant of Rudebusch and Svensson [14]) and an estimated forward-looking model (a three-mode variant of Lindé [11]). Svensson and Williams [16] also extend the MJLQ framework to the more realistic case of unobservable modes, although without introducing learning and inference about the probability distribution of modes. Svensson and Williams [17] focus on learning and experimentation in the MJLQ framework.

The paper is organized as follows: Section 2 presents the MJLQ framework and summarizes our earlier work. Section 3 presents our analysis of learning and experimentation in a simple benchmark New Keynesian model, whereas section 4 presents our analysis in an estimated empirical New Keynesian model. Section 5 presents some conclusions and suggestions for further work.

2 MJLQ Analysis of Optimal Policy

This section summarizes our earlier work, Svensson and Williams [16] and [17].

2.1 An MJLQ model

We consider an MJLQ model of an economy with forward-looking variables. The economy has a private sector and a policymaker. We let X_t denote an n_X -vector of predetermined variables in period t , x_t an n_x -vector of forward-looking variables, and i_t an n_i -vector of (policymaker) instruments (control variables).⁵ We let model uncertainty be represented by n_j possible modes and let $j_t \in N_j \equiv \{1, 2, \dots, n_j\}$ denote the mode in period t . The model of the economy can then be written

$$X_{t+1} = A_{11j_{t+1}}X_t + A_{12j_{t+1}}x_t + B_{1j_{t+1}}i_t + C_{1j_{t+1}}\varepsilon_{t+1}, \quad (2.1)$$

$$E_t H_{j_{t+1}} x_{t+1} = A_{21j_t}X_t + A_{22j_t}x_t + B_{2j_t}i_t + C_{2j_t}\varepsilon_t, \quad (2.2)$$

where ε_t is a multivariate normally distributed random i.i.d. n_ε -vector of shocks with mean zero and contemporaneous covariance matrix I_{n_ε} . The matrices A_{11j} , A_{12j} , ..., C_{2j} have the appropriate dimensions and depend on the mode j . As a structural model here is simply a collection of matrices, each mode can represent a different model of the economy. Thus, uncertainty about the prevailing mode *is* model uncertainty.⁶

Note that the matrices on the right side of (2.1) depend on the mode j_{t+1} in period $t + 1$, whereas the matrices on the right side of (2.2) depend on the mode j_t in period t . Equation (2.1) then determines the predetermined variables in period $t + 1$ as a function of the mode and shocks in period $t + 1$ and the predetermined variables, forward-looking variables, and instruments in period t . Equation (2.2) determines the forward-looking variables in period t as a function of the mode and shocks in period t , the expectations in period t of next period's mode and forward-looking variables,

⁵ The first component of X_t may be unity, in order to allow for mode-dependent intercepts in the model equations.

⁶ See also Svensson and Williams [16], where we show how many different types of uncertainty can be mapped into our MJLQ framework.

and the predetermined variables and instruments in period t . The matrix A_{22j} is non-singular for each $j \in N_j$.

The mode j_t follows a Markov process with the transition matrix $P \equiv [P_{jk}]$.⁷ The shocks ε_t are mean zero and i.i.d. with probability density φ , and without loss of generality we assume that ε_t is independent of j_t .⁸ We also assume that $C_{1j}\varepsilon_t$ and $C_{2k}\varepsilon_t$ are independent for all $j, k \in N_j$. These shocks, along with the modes, are the driving forces in the model. They are not directly observed. For technical reasons, it is convenient but not necessary that they are independent. We let $p_t = (p_{1t}, \dots, p_{n_j t})'$ denote the true probability distribution of j_t in period t . We let $p_{t+\tau|t}$ denote the policymaker's and private sector's estimate in the beginning of period t of the probability distribution in period $t + \tau$. The *prediction* equation for the probability distribution is

$$p_{t+1|t} = P' p_{t|t}. \quad (2.3)$$

We let the operator $E_t[\cdot]$ in the expression $E_t H_{j_{t+1}} x_{t+1}$ on the left side of (2.2) denote expectations in period t conditional on policymaker and private-sector information in the beginning of period t , including X_t , i_t , and $p_{t|t}$, but excluding j_t and ε_t . Thus, the maintained assumption is symmetric information between the policymaker and the (aggregate) private sector. Since forward-looking variables will be allowed to depend on j_t , *parts* of the private sector, but not the *aggregate* private sector, may be able to observe j_t and parts of ε_t . Note that although we focus on the determination of the optimal policy instrument i_t , our results also show how private sector choices as embodied in x_t are affected by uncertainty and learning. The precise informational assumptions and the determination of $p_{t|t}$ will be specified below.

We let the policymaker's intertemporal loss function in period t be

$$E_t \sum_{\tau=0}^{\infty} \delta^\tau L(X_{t+\tau}, x_{t+\tau}, i_{t+\tau}, j_{t+\tau}) \quad (2.4)$$

where δ is a discount factor satisfying $0 < \delta < 1$, and the period loss, $L(X_t, x_t, i_t, j_t)$, satisfies

$$L(X_t, x_t, i_t, j_t) \equiv \begin{bmatrix} X_t \\ x_t \\ i_t \end{bmatrix}' W_{j_t} \begin{bmatrix} X_t \\ x_t \\ i_t \end{bmatrix}, \quad (2.5)$$

where the matrix W_j ($j \in N_j$) is positive semidefinite. We assume that the policymaker optimizes under commitment in a timeless perspective. As explained below, we will then add the term

$$\Xi_{t-1} \frac{1}{\delta} E_t H_{j_t} x_t \quad (2.6)$$

⁷ Obvious special cases are $P = I_{n_j}$, when the modes are completely persistent, and $P_j = \bar{p}'$ ($j \in N_j$), when the modes are serially i.i.d. with probability distribution \bar{p} .

⁸ Because mode-dependent intercepts (as well as mode-dependent standard deviations) are allowed in the model, we can still incorporate additive mode-dependent shocks.

to the intertemporal loss function in period t . As we shall see below, the n_x -vector Ξ_{t-1} is the vector of Lagrange multipliers for equation (2.2) from the optimization problem in period $t - 1$. For the special case when there are no forward-looking variables ($n_x = 0$), the model consists of (2.1) only, without the term $A_{12j_{t+1}}x_t$; the period loss function depends on X_t , i_t , and j_t only; and there is no role for the Lagrange multipliers Ξ_{t-1} or the term (2.6).

2.2 Approximate MJLQ models

While in this paper we start with an MJLQ model, it is natural to ask where such a model comes from, as usual formulations of economic models are not of this type. However the same type of approximation methods that are widely used to convert nonlinear models into their linear counterparts can also convert nonlinear models into MJLQ models. We analyze this issue in Svensson and Williams [16], and present an illustration as well. Here we briefly discuss the main ideas. Rather than analyzing local deviations from a single steady state as in conventional linearizations, for an MJLQ approximation we analyze the local deviations from (potentially) separate, mode-dependent steady states. Standard linearizations are justified as asymptotically valid for small shocks, as an increasing time is spent in the vicinity of the steady state. Our MJLQ approximations are asymptotically valid for small shocks and persistent modes, as an increasing time is spent in the vicinity of each mode-dependent steady state. Thus, for slowly-varying Markov chains, our MJLQ provide accurate approximations of nonlinear models with Markov switching.

2.3 Types of optimal policies

We will distinguish three cases: (1) Optimal policy when there is no learning (NL), (2) Adaptive optimal policy (AOP), and (3) Bayesian optimal policy (BOP). By NL, we refer to a situation when the policymaker and the aggregate private sector have a probability distribution $p_{t|t}$ over the modes in period t and updates the probability distribution in future periods using the transition matrix only, so the *updating* equation is

$$p_{t+1|t+1} = P^t p_{t|t}. \quad (2.7)$$

That is, the policymaker and the private sector do not use observations of the variables in the economy to update the probability distribution. The policymaker then determines optimal policy in period t conditional on $p_{t|t}$ and (2.7). This is a variant of a case examined in Svensson and Williams [16].

By AOP, we refer to a situation when the policymaker in period t determines optimal policy as in the NL case, but then uses observations of the realization of the variables in the economy to update its probability distribution according to Bayes Theorem. In this case, the instruments will generally have an effect on the updating of future probability distributions, and through this channel separately affect the intertemporal loss. However, the policymaker does not exploit that channel in determining optimal policy. That is, the policymaker does not do any conscious experimentation. By BOP, we refer to a situation when the policymaker acknowledges that the current instruments will affect future inference and updating of the probability distribution, and calculates optimal policy taking this separate channel into account. Therefore, BOP includes optimal experimentation, where for instance the policymaker may pursue policy that increases losses in the short run but improves the inference of the probability distribution and therefore lowers losses in the longer run.

2.4 Optimal policy with no learning

We first consider the NL case. Svensson and Williams [16] derive the equilibrium under commitment in a timeless perspective for the case when X_t , x_t , and i_t are observable in period t , j_t is unobservable, and the updating equation for $p_{t|t}$ is given by (2.7). Observations of X_t , x_t , and i_t are then not used to update $p_{t|t}$.

It will be useful to replace equation (2.2) by the two equivalent equations,

$$E_t H_{j_{t+1}} x_{t+1} = z_t, \tag{2.8}$$

$$0 = A_{21j_t} X_t + A_{22j_t} x_t - z_t + B_{2j_t} i_t + C_{2j_t} \varepsilon_t, \tag{2.9}$$

where we introduce the n_x -vector of additional forward-looking variables, z_t . Introducing this vector is a practical way of keeping track of the expectations term on the left side of (2.2). Furthermore, it will be practical to use (2.9) and solve x_t as a function of X_t , z_t , i_t , j_t , and ε_t

$$x_t = \tilde{x}(X_t, z_t, i_t, j_t, \varepsilon_t) \equiv A_{22j_t}^{-1} (z_t - A_{21j_t} X_t - B_{2j_t} i_t - C_{2j_t} \varepsilon_t). \tag{2.10}$$

We note that, for given j_t , this function is linear in X_t , z_t , i_t , and ε_t .

In order to solve for the optimal decisions, we use the recursive saddlepoint method (see Marcet and Marimon [12], Svensson and Williams [16], and Svensson [15] for details of the recursive saddlepoint method). Thus, we introduce Lagrange multipliers for each forward-looking equation, the lagged values of which become state variables and reflect costs of commitment, while the current

values become control variables. The dual period loss function can be written

$$\mathbf{E}_t \tilde{L}(\tilde{X}_t, z_t, i_t, \gamma_t, j_t, \varepsilon_t) \equiv \sum_j p_{jt|t} \int \tilde{L}(\tilde{X}_t, z_t, i_t, \gamma_t, j, \varepsilon_t) \varphi(\varepsilon_t) d\varepsilon_t,$$

where $\tilde{X}_t \equiv (X'_t, \Xi'_{t-1})'$ is the $(n_X + n_x)$ -vector of extended predetermined variables (that is, including the n_x -vector Ξ_{t-1}), γ_t is an n_x -vector of Lagrange multipliers, and $\varphi(\cdot)$ denotes a generic probability density function (for ε_t , the standard normal density function), and where

$$\tilde{L}(\tilde{X}_t, z_t, i_t, \gamma_t, j_t, \varepsilon_t) \equiv L[X_t, \tilde{x}(X_t, z_t, i_t, j_t, \varepsilon_t), i_t, j_t] - \gamma'_t z_t + \Xi'_{t-1} \frac{1}{\delta} H_{j_t} \tilde{x}(X_t, z_t, i_t, j_t, \varepsilon_t). \quad (2.11)$$

As discussed in Svensson and Williams [16], the failure of the law of iterated expectations leads us to introduce the collection of value functions $\hat{V}(s_t, j)$ which condition on the mode, while the value function $\tilde{V}(s_t)$ averages over these and represents the solution of the dual optimization problem. The somewhat unusual Bellman equation for the dual problem can be written

$$\begin{aligned} \tilde{V}(s_t) &\equiv \mathbf{E}_t \hat{V}(s_t, j_t) \equiv \sum_j p_{jt|t} \hat{V}(s_t, j) \\ &= \max_{\gamma_t} \min_{(z_t, i_t)} \mathbf{E}_t \{ \tilde{L}(\tilde{X}_t, z_t, i_t, \gamma_t, j_t, \varepsilon_t) + \delta \hat{V}[g(s_t, z_t, i_t, \gamma_t, j_t, \varepsilon_t, j_{t+1}, \varepsilon_{t+1}), j_{t+1}] \} \\ &\equiv \max_{\gamma_t} \min_{(z_t, i_t)} \sum_j p_{jt|t} \int \left[\tilde{L}(\tilde{X}_t, z_t, i_t, \gamma_t, j, \varepsilon_t) \right. \\ &\quad \left. + \delta \sum_k P_{jk} \hat{V}[g(s_t, z_t, i_t, \gamma_t, j, \varepsilon_t, k, \varepsilon_{t+1}), k] \right] \varphi(\varepsilon_t) \varphi(\varepsilon_{t+1}) d\varepsilon_t d\varepsilon_{t+1}. \end{aligned} \quad (2.12)$$

where $s_t \equiv (\tilde{X}'_t, p'_{t|t})'$ denotes the *perceived* state of the economy (it includes the perceived probability distribution, $p_{t|t}$, but not the true mode) and (s_t, j_t) denotes the *true* state of the economy (it includes the true mode of the economy). As we discuss in more detail below, it is necessary to include the mode j_t in the state vector because the beliefs do not satisfy the law of iterated expectations. In the BOP case beliefs do satisfy this property, so the state vector is simply s_t . Also note that in the Bellman equation we require that all the choice variables respect the information constraints, and thus depend on the perceived state s_t but not the mode j directly.

The optimization is subject to the transition equation for X_t ,

$$X_{t+1} = A_{11j_{t+1}} X_t + A_{12j_{t+1}} \tilde{x}(X_t, z_t, i_t, j_t, \varepsilon_t) + B_{1j_{t+1}} i_t + C_{1j_{t+1}} \varepsilon_{t+1}, \quad (2.13)$$

where we have substituted $\tilde{x}(X_t, z_t, i_t, j_t, \varepsilon_t)$ for x_t ; the new dual transition equation for Ξ_t ,

$$\Xi_t = \gamma_t, \quad (2.14)$$

and the transition equation (2.7) for $p_{t|t}$. Combining equations, we have the transition for s_t ,

$$\begin{aligned} s_{t+1} &\equiv \begin{bmatrix} X_{t+1} \\ \Xi_t \\ p_{t+1|t+1} \end{bmatrix} = g(s_t, z_t, i_t, \gamma_t, j_t, \varepsilon_t, j_{t+1}, \varepsilon_{t+1}) \\ &\equiv \begin{bmatrix} A_{11j_{t+1}}X_t + A_{12j_{t+1}}\tilde{x}(X_t, z_t, i_t, j, \varepsilon_t) + B_{1j_{t+1}}i_t + C_{1j_{t+1}}\varepsilon_{t+1} \\ \gamma_t \\ P'p_{t|t} \end{bmatrix}. \end{aligned} \quad (2.15)$$

It is straightforward to see that the solution of the dual optimization problem (2.12) is linear in \tilde{X}_t for given $p_{t|t}, j_t$,

$$\begin{bmatrix} z_t \\ i_t \\ \gamma_t \end{bmatrix} = \begin{bmatrix} z(s_t) \\ i(s_t) \\ \gamma(s_t) \end{bmatrix} = F(p_{t|t})\tilde{X}_t \equiv \begin{bmatrix} F_z(p_{t|t}) \\ F_i(p_{t|t}) \\ F_\gamma(p_{t|t}) \end{bmatrix} \tilde{X}_t, \quad (2.16)$$

$$x_t = x(s_t, j_t, \varepsilon_t) \equiv \tilde{x}(X_t, z(s_t), i(s_t), j_t, \varepsilon_t) \equiv F_{x\tilde{X}}(p_{t|t}, j_t)\tilde{X}_t + F_{x\varepsilon}(p_{t|t}, j_t)\varepsilon_t. \quad (2.17)$$

This solution is also the solution to the original primal optimization problem. We note that x_t is linear in ε_t for given $p_{t|t}$ and j_t . The equilibrium transition equation is then given by

$$s_{t+1} = \hat{g}(s_t, j_t, \varepsilon_t, j_{t+1}, \varepsilon_{t+1}) \equiv g[s_t, z(s_t), i(s_t), \gamma(s_t), j_t, \varepsilon_t, j_{t+1}, \varepsilon_{t+1}]. \quad (2.18)$$

As can be easily verified, the (unconditional) dual value function $\tilde{V}(s_t)$ is quadratic in \tilde{X}_t for given $p_{t|t}$, taking the form

$$\tilde{V}(s_t) \equiv \tilde{X}_t' \tilde{V}_{\tilde{X}\tilde{X}}(p_{t|t}) \tilde{X}_t + w(p_{t|t}).$$

The conditional dual value function $\hat{V}(s_t, j_t)$ gives the dual intertemporal loss conditional on the true state of the economy, (s_t, j_t) . It follows that this function satisfies

$$\hat{V}(s_t, j) \equiv \int \left[\tilde{L}(\tilde{X}_t, z(s_t), i(s_t), \gamma(s_t), j, \varepsilon_t) + \delta \sum_k P_{jk} \hat{V}[\hat{g}(s_t, j, \varepsilon_t, k, \varepsilon_{t+1}), k] \right] \varphi(\varepsilon_t) \varphi(\varepsilon_{t+1}) d\varepsilon_t d\varepsilon_{t+1} \quad (j \in N_j).$$

The function $\hat{V}(s_t, j_t)$ is also quadratic in \tilde{X}_t for given $p_{t|t}$ and j_t ,

$$\hat{V}(s_t, j_t) \equiv \tilde{X}_t' \hat{V}_{\tilde{X}\tilde{X}}(p_{t|t}, j_t) \tilde{X}_t + \hat{w}(p_{t|t}, j_t).$$

It follows that we have

$$\tilde{V}_{\tilde{X}\tilde{X}}(p_{t|t}) \equiv \sum_j p_{j|t} \hat{V}_{\tilde{X}\tilde{X}}(p_{t|t}, j), \quad w(p_{t|t}) \equiv \sum_j p_{j|t} \hat{w}(p_{t|t}, j).$$

Although we find the optimal policies from the dual problem, in order to measure true expected losses we are interested in the value function for the primal problem (with the original, unmodified

loss function). This value function, with the period loss function $E_t L(X_t, x_t, i_t, j_t)$ rather than $E_t \tilde{L}(\tilde{X}_t, z_t, i_t, \gamma_t, j_t, \varepsilon_t)$, satisfies

$$\begin{aligned} V(s_t) &\equiv \tilde{V}(s_t) - \Xi'_{t-1} \frac{1}{\delta} \sum_j p_{jt|t} H_j \int x(s_t, j, \varepsilon_t) \varphi(\varepsilon_t) d\varepsilon_t \\ &= \tilde{V}(s_t) - \Xi'_{t-1} \frac{1}{\delta} \sum_j p_{jt|t} H_j x(s_t, j, 0) \end{aligned} \quad (2.19)$$

(where the second equality follows since $x(s_t, j_t, \varepsilon_t)$ is linear in ε_t for given s_t and j_t). It is quadratic in \tilde{X}_t for given $p_{t|t}$,

$$V(s_t) \equiv \tilde{X}'_t V_{\tilde{X}\tilde{X}}(p_{t|t}) \tilde{X}_t + w(p_{t|t})$$

(the scalar $w(p_{t|t})$ in the primal value function is obviously identical to that in the dual value function). This is the value function conditional on \tilde{X}_t and $p_{t|t}$ after X_t has been observed but before x_t has been observed, taking into account that j_t and ε_t are not observed. Hence, the second term on the right side of (2.19) contains the expectation of $H_{j_t} x_t$ conditional on that information.⁹

Svensson and Williams [16] and [17] present algorithms to compute the solution and the primal and dual value functions for the no-learning case. For future reference, we note that the value function for the primal problem also satisfies

$$V(s_t) \equiv \sum_j p_{jt|t} \check{V}(s_t, j),$$

where the conditional value function, $\check{V}(s_t, j_t)$, satisfies

$$\check{V}(s_t, j) = \int \left\{ \begin{array}{l} L[X_t, x(s_t, j, \varepsilon_t), i(s_t), j] \\ + \delta \sum_k P_{jk} \check{V}[\hat{g}(s_t, j, \varepsilon_t, k, \varepsilon_{t+1}), k] \end{array} \right\} \varphi(\varepsilon_t) \varphi(\varepsilon_{t+1}) d\varepsilon_t d\varepsilon_{t+1} \quad (j \in N_j). \quad (2.20)$$

2.5 Adaptive optimal policy

Consider now the case of adaptive optimal policy, where the policymaker uses the same policy function as in the no-learning case, but each period updates the probabilities that this policy is conditioned on. This case is thus simple to implement recursively, as we have already discussed how to solve for the optimal decisions and below we show how to update probabilities. However, the ex-ante evaluation of expected loss is more complex, as we show below. In particular, we assume that $C_{2j_t} \neq 0$ and that both ε_t and j_t are unobservable. The estimate $p_{t|t}$ is the result of Bayesian updating, using all information available, but the optimal policy in period t is computed under

⁹ To be precise, the observation of X_t , which depends on $C_{1j_t} \varepsilon_t$, allows some inference of ε_t , $\varepsilon_{t|t}$. x_t will depend on j_t and on ε_t , but on ε_t only through $C_{2j_t} \varepsilon_t$. By assumption $C_{1j} \varepsilon_t$ and $C_{2k} \varepsilon_t$ are independent. Hence, any observation of X_t and $C_{1j} \varepsilon_t$ does not convey any information about $C_{2j} \varepsilon_t$, so $E_t C_{2j_t} \varepsilon_t = 0$.

the perceived updating equation (2.7). That is, the fact that the policy choice will affect future $p_{t+\tau|t+\tau}$ and that future expected loss will change when $p_{t+\tau|t+\tau}$ changes is disregarded. Under the assumption that the expectations on the left side of (2.2) are conditional on (2.7), the variables z_t , i_t , γ_t , and x_t in period t are still determined by (2.16) and (2.17).

In order to determine the updating equation for $p_{t|t}$, we specify an explicit sequence of information revelation as follows, in no less than nine steps. The timing assumptions are necessary in order to spell out the appropriate conditioning for decisions and updating of beliefs.

First, the policymaker and the private sector enters period t with the prior $p_{t|t-1}$. They know X_{t-1} , $x_{t-1} = x(s_{t-1}, j_{t-1}, \varepsilon_{t-1})$, $z_{t-1} = z(s_{t-1})$, $i_{t-1} = i(s_{t-1})$, and $\Xi_{t-1} = \gamma(s_{t-1})$ from the previous period.

Second, in the beginning of period t , the mode j_t and the vector of shocks ε_t are realized. Then the vector of predetermined variables X_t is realized according to (2.1).

Third, the policymaker and the private sector observe X_t . They then know $\tilde{X}_t \equiv (X_t', \Xi_{t-1}')'$. They do not observe j_t or ε_t .

Fourth, the policymaker and the private sector update the prior $p_{t|t-1}$ to the posterior $p_{t|t}$ according to Bayes Theorem and the updating equation

$$p_{j_t|t} = \frac{\varphi(X_t|j_t = j, X_{t-1}, x_{t-1}, i_{t-1}, p_{t|t-1})}{\varphi(X_t|X_{t-1}, x_{t-1}, i_{t-1}, p_{t|t-1})} p_{j_t|t-1} \quad (j \in N_j), \quad (2.21)$$

where again $\varphi(\cdot)$ denotes a generic density function.¹⁰ Then the policymaker and the private sector know $s_t \equiv (\tilde{X}_t', p_{t|t}')'$.

Fifth, the policymaker solves the dual optimization problem, determines $i_t = i(s_t)$, and implements/announces the instrument setting i_t .

Sixth, the private-sector (and policymaker) expectations,

$$z_t = \mathbb{E}_t H_{j_{t+1}} x_{t+1} \equiv \mathbb{E}[H_{j_{t+1}} x_{t+1} | s_t],$$

are formed. In equilibrium, these expectations will be determined by (2.16). In order to understand their determination better, we look at this in some detail.

These expectations are by assumption formed before x_t is observed. The private sector and the policymaker know that x_t will in equilibrium be determined in the next step according to (2.17).

¹⁰ The policymaker and private sector can also estimate the shocks $\varepsilon_{t|t}$ as $\varepsilon_{t|t} = \sum_j p_{j_t|t} \varepsilon_{j_t|t}$, where $\varepsilon_{j_t|t} \equiv X_t - A_{11j} X_{t-1} - A_{12j} x_{t-1} - B_{1j} i_{t-1}$ ($j \in N_j$). However, because of the assumed independence of $C_{1j} \varepsilon_t$ and $C_{2k} \varepsilon_t$, $j, k \in N_j$, we do not need to keep track of $\varepsilon_{j_t|t}$.

Hence, they can form expectations of the soon-to-be determined x_t conditional on $j_t = j$,¹¹

$$x_{jt|t} = x(s_t, j, 0). \quad (2.22)$$

The private sector and the policymaker can also infer Ξ_t from

$$\Xi_t = \gamma(s_t). \quad (2.23)$$

This allows the private sector and the policymaker to form the expectations

$$z_t = z(s_t) = \mathbb{E}_t[H_{j_{t+1}}x_{t+1} | s_t] = \sum_{j,k} P_{jk}p_{jt|t}H_kx_{k,t+1|jt}, \quad (2.24)$$

where

$$\begin{aligned} x_{k,t+1|jt} &= \int x \left(\begin{bmatrix} A_{11k}X_t + A_{12k}x(s_t, j, \varepsilon_t) + B_{1k}i(s_t) \\ \Xi_t \\ P'p_{t|t} \end{bmatrix}, k, \varepsilon_{t+1} \right) \varphi(\varepsilon_t)\varphi(\varepsilon_{t+1})d\varepsilon_t d\varepsilon_{t+1} \\ &= x \left(\begin{bmatrix} A_{11k}X_t + A_{12k}x(s_t, j, 0) + B_{1k}i(s_t) \\ \Xi_t \\ P'p_{t|t} \end{bmatrix}, k, 0 \right), \end{aligned}$$

where we have exploited the linearity of $x_t = x(s_t, j_t, \varepsilon_t)$ and $x_{t+1} = x(s_{t+1}, j_{t+1}, \varepsilon_{t+1})$ in ε_t and ε_{t+1} . Note that z_t is, under AOP, formed conditional on the belief that the probability distribution in period $t + 1$ will be given by $p_{t+1|t+1} = P'p_{t|t}$, not by the true updating equation that we are about to specify.

Seventh, after the expectations z_t have been formed, x_t is determined as a function of X_t , z_t , i_t , j_t , and ε_t by (2.10).

Eighth, the policymaker and the private sector then use the observed x_t to update $p_{t|t}$ to the new posterior $p_{t|t}^+$ according to Bayes Theorem, via the updating equation

$$p_{jt|t}^+ = \frac{\varphi(x_t|j_t = j, X_t, z_t, i_t, p_{t|t})}{\varphi(x_t|X_t, z_t, i_t, p_{t|t})}p_{jt|t} \quad (j \in N_j). \quad (2.25)$$

Ninth, the policymaker and the private sector then leave period t and enter period $t + 1$ with the prior $p_{t+1|t}$ given by the prediction equation

$$p_{t+1|t} = P'p_{t|t}^+. \quad (2.26)$$

In the beginning of period $t + 1$, the mode j_{t+1} and the vector of shocks ε_{t+1} are realized, and X_{t+1} is determined by (2.1) and observed by the policymaker and private sector. The sequence of

¹¹ Note that 0 instead of $\varepsilon_{jt|t}$ enters above. This is because the inference $\varepsilon_{jt|t}$ above is inference about $C_{1j}\varepsilon_t$, whereas x_t depends on ε_t through $C_{2j}\varepsilon_t$. Since we assume that $C_{1j}\varepsilon_t$ and $C_{2j}\varepsilon_t$ are independent, there is no inference of $C_{2j}\varepsilon_t$ from observing X_t . Hence, $\mathbb{E}_t C_{2j_t}\varepsilon_t \equiv 0$. Because of the linearity of x_t in ε_t , the integration of x_t over ε_t results in $x(s_t, j_t, 0_t)$.

the nine steps above then repeats itself. For more detail on the explicit densities in the updating equations (2.21) and (2.25) see Svensson and Williams [17].

The transition equation for $p_{t+1|t+1}$ can be written

$$p_{t+1|t+1} = Q(s_t, z_t, i_t, j_t, \varepsilon_t, j_{t+1}, \varepsilon_{t+1}), \quad (2.27)$$

where $Q(s_t, z_t, i_t, j_t, \varepsilon_t, j_{t+1}, \varepsilon_{t+1})$ is defined by the combination of (2.21) for period $t+1$ with (2.13) and (2.26). The equilibrium transition equation for the full state vector is then given by

$$\begin{aligned} s_{t+1} &\equiv \begin{bmatrix} X_{t+1} \\ \Xi_t \\ p_{t+1|t+1} \end{bmatrix} = \bar{g}(s_t, j_t, \varepsilon_t, j_{t+1}, \varepsilon_{t+1}) \\ &\equiv \begin{bmatrix} A_{11j_{t+1}} X_t + A_{12j_{t+1}} x(s_t, j_t, \varepsilon_t) + B_{1j_{t+1}} i(s_t) + C_{1j_{t+1}} \varepsilon_{t+1} \\ \gamma(s_t) \\ Q(s_t, z(s_t), i(s_t), j_t, \varepsilon_t, j_{t+1}, \varepsilon_{t+1}) \end{bmatrix}, \end{aligned} \quad (2.28)$$

where the third row is given by the true updating equation (2.27) together with the policy function (2.16). Thus, we note that, in this AOP case, there is a distinction between the “perceived” transition and equilibrium transition equations, (2.15) and (2.18), which in the bottom block include the perceived updating equation (2.7), and the “true” equilibrium transition equation, (2.28), which replaces the perceived updating equation (2.7) with the true updating equation (2.27).

Note that $V(s_t)$ in (2.19), which is subject to the perceived transition equation, (2.15), does not give the true (unconditional) value function for the AOP case. This is instead given by

$$\bar{V}(s_t) \equiv \sum_j p_{jt|t} \check{V}(s_t, j),$$

where the true conditional value function, $\check{V}(s_t, j)$, satisfies

$$\check{V}(s_t, j) = \int \left\{ \begin{array}{l} L[X_t, x(s_t, j, \varepsilon_t), i(s_t), j] \\ + \delta \sum_k P_{jk} \check{V}[\bar{g}(s_t, j, \varepsilon_t, k, \varepsilon_{t+1}), k] \end{array} \right\} \varphi(\varepsilon_t) \varphi(\varepsilon_{t+1}) d\varepsilon_t d\varepsilon_{t+1} \quad (j \in N_j). \quad (2.29)$$

That is, the true value function $\bar{V}(s_t)$ takes into account the true updating equation for $p_{t|t}$, (2.27), whereas the optimal policy, (2.16), and the perceived value function, $V(s_t)$ in (2.19), are conditional on the perceived updating equation (2.7) and thereby the perceived transition equation (2.15). Note also that $\bar{V}(s_t)$ is the value function after \tilde{X}_t has been observed but before x_t is observed, so it is conditional on $p_{t|t}$ rather than $p_{t|t}^+$. Since the full transition equation (2.28) is no longer linear due to the belief updating (2.27), the true value function $\bar{V}(s_t)$ is no longer quadratic in \tilde{X}_t for given $p_{t|t}$. Thus, more complex numerical methods are required to evaluate losses in the AOP case, although policy is still determined simply as in the NL case.

As we discuss in Svensson and Williams [17], the difference between the true updating equation for $p_{t+1|t+1}$, (2.27), and the perceived updating equation (2.7) is that, in the true updating equation, $p_{t+1|t+1}$ becomes a random variable from the point of view of period t , with mean equal to $p_{t+1|t}$. This is because $p_{t+1|t+1}$ depends on the realization of j_{t+1} and ε_{t+1} . Thus Bayesian updating induces a mean-preserving spread over beliefs, which in turn sheds light on the gains from learning. If the conditional value function $\check{V}(s_t, j_t)$ under NL is concave in $p_{t|t}$ for given \check{X}_t and j_t , then by Jensen's inequality the true expected future loss under AOP will be lower than the true expected future loss under NL. That is, the concavity of the value function in beliefs means that learning leads to lower losses. While it likely that \check{V} is indeed concave, as we show in applications, it need not be globally so and thus learning need not always reduce losses. In some cases the losses incurred by increased variability of beliefs may offset the expected precision gains. Furthermore, under BOP, it may be possible to adjust policy so as to further increase the variance of $p_{t|t}$, that is, achieve a mean-preserving spread which might further reduce the expected future loss.¹² This amounts to optimal experimentation.

2.6 Bayesian optimal policy

Finally, we consider the BOP case, when optimal policy is determined while taking the updating equation (2.27) into account. That is, we now allow the policymaker to choose i_t taking into account that his actions will affect $p_{t+1|t+1}$, which in turn will affect future expected losses. In particular, experimentation is allowed and is optimally chosen. For the BOP case, there is hence no distinction between the “perceived” and “true” transition equation.

The transition equation for the BOP case is:

$$\begin{aligned}
s_{t+1} &\equiv \begin{bmatrix} X_{t+1} \\ \Xi_t \\ p_{t+1|t+1} \end{bmatrix} = g(s_t, z_t, i_t, \gamma_t, j_t, \varepsilon_t, j_{t+1}, \varepsilon_{t+1}) \\
&\equiv \begin{bmatrix} A_{11j_{t+1}} X_t + A_{12j_{t+1}} \tilde{x}(s_t, z_t, i_t, j_t, \varepsilon_t) + B_{1j_{t+1}} i_t + C_{1j_{t+1}} \varepsilon_{t+1} \\ \gamma_t \\ Q(s_t, z_t, i_t, j_t, \varepsilon_t, j_{t+1}, \varepsilon_{t+1}) \end{bmatrix}. \quad (2.30)
\end{aligned}$$

Then the dual optimization problem can be written as (2.12) subject to the above transition equation (2.30). However, in the Bayesian case, matters simplify somewhat, as we do not need to compute the conditional value functions $\hat{V}(s_t, j_t)$, which we recall were required due to the failure of the law of iterated expectations in the AOP case. We note now that the second term on the

¹² Kiefer [10] examines the properties of a value function, including concavity, under Bayesian learning for a simpler model without forward looking variables.

right side of (2.12) can be written as

$$\mathbb{E}_t \hat{V}(s_{t+1}, j_{t+1}) \equiv \mathbb{E} \left[\hat{V}(s_{t+1}, j_{t+1}) \mid s_t \right].$$

Since, in the Bayesian case, the beliefs do satisfy the law of iterated expectations, this is then the same as

$$\mathbb{E} \left[\hat{V}(s_{t+1}, j_{t+1}) \mid s_t \right] = \mathbb{E} \left[\tilde{V}(s_{t+1}) \mid s_t \right].$$

See Svensson and Williams [17] for a proof.

Thus, the dual Bellman equation for the Bayesian optimal policy is

$$\begin{aligned} \tilde{V}(s_t) &= \max_{\gamma_t} \min_{(z_t, i_t)} \mathbb{E}_t \{ \tilde{L}(\tilde{X}_t, z_t, i_t, \gamma_t, j_t, \varepsilon_t) + \delta \tilde{V}[g(s_t, z_t, i_t, \gamma_t, j_t, \varepsilon_t, j_{t+1}, \varepsilon_{t+1})] \} \\ &\equiv \max_{\gamma_t} \min_{(z_t, i_t)} \sum_j p_{jt|t} \int \left[\begin{array}{l} \tilde{L}(\tilde{X}_t, z_t, i_t, \gamma_t, j, \varepsilon_t) \\ + \delta \sum_k P_{jk} \tilde{V}[g(s_t, z_t, i_t, \gamma_t, j, \varepsilon_t, k, \varepsilon_{t+1})] \end{array} \right] \varphi(\varepsilon_t) \varphi(\varepsilon_{t+1}) d\varepsilon_t d\varepsilon_{t+1}, \end{aligned} \quad (2.31)$$

where the transition equation is given by (2.30).

The solution to the optimization problem can be written

$$\tilde{i}_t \equiv \begin{bmatrix} z_t \\ i_t \\ \gamma_t \end{bmatrix} = \tilde{i}(s_t) \equiv \begin{bmatrix} z(s_t) \\ i(s_t) \\ \gamma(s_t) \end{bmatrix} = F(\tilde{X}_t, p_{t|t}) \equiv \begin{bmatrix} F_z(\tilde{X}_t, p_{t|t}) \\ F_i(\tilde{X}_t, p_{t|t}) \\ F_\gamma(\tilde{X}_t, p_{t|t}) \end{bmatrix}, \quad (2.32)$$

$$x_t = x(s_t, j_t, \varepsilon_t) \equiv \tilde{x}(X_t, z(s_t), i(s_t), j_t, \varepsilon_t) \equiv F_x(\tilde{X}_t, p_{t|t}, j_t, \varepsilon_t). \quad (2.33)$$

Because of the nonlinearity of (2.27) and (2.30), the solution is no longer linear in \tilde{X}_t for given $p_{t|t}$. The dual value function, $\tilde{V}(s_t)$, is no longer quadratic in \tilde{X}_t for given $p_{t|t}$. The value function of the primal problem, $V(s_t)$, is given by, equivalently, (2.19), (2.29) (with the equilibrium transition equation (2.28) with the solution (2.32)), or

$$V(s_t) = \sum_j p_{jt|t} \int \left\{ \begin{array}{l} L[X_t, x(s_t, j, \varepsilon_t), i(s_t), j] \\ + \delta \sum_k P_{jk} V[\bar{g}(s_t, j, \varepsilon_t, k, \varepsilon_{t+1})] \end{array} \right\} \varphi(\varepsilon_t) \varphi(\varepsilon_{t+1}) d\varepsilon_t d\varepsilon_{t+1}. \quad (2.34)$$

It is also no longer quadratic in \tilde{X}_t for given $p_{t|t}$. Thus, more complex and detailed numerical methods are necessary in this case to find the optimal policy and the value function. Therefore little can be said in general about the solution of the problem. Nonetheless, in numerical analysis it is very useful to have a good starting guess at a solution, which in our case comes from the AOP case. In our examples below we explain in more detail how the BOP and AOP cases differ, and what drives the differences.

3 Learning and experimentation in a simple New Keynesian model

3.1 The model

We consider the benchmark standard New Keynesian model, consisting of a New Keynesian Phillips curve and a consumption Euler equation (see Woodford [21] for an exposition):

$$\pi_t = \delta E_t \pi_{t+1} + \gamma_{j_t} y_t + c_\pi \varepsilon_{\pi t}, \quad (3.1)$$

$$y_t = E_t y_{t+1} - \sigma_{j_t} (i_t - E_t \pi_{t+1}) + c_y \varepsilon_{y_t} + c_g g_t, \quad (3.2)$$

$$g_{t+1} = \rho g_t + \varepsilon_{g,t+1}. \quad (3.3)$$

Here π_t is the inflation rate, y_t is the output gap, δ is the subjective discount factor (as above), γ_{j_t} is a composite parameter reflecting the elasticity of demand and frequency of price adjustment, and σ_{j_t} is the intertemporal elasticity of substitution. There are three shocks in the model, two unobservable shocks $\varepsilon_{\pi t}$ and ε_{y_t} , which are independent standard normal random variables, and the observable serially correlated shock g_t . This last shock is interpretable as a “demand” shock either coming from variation in preferences, government spending, or the underlying efficient level of output. Woodford [21] combines and renormalizes these shocks into a composite shock representing variation in the natural rate of interest.

In the standard formulations of this model, the shocks are observable and policy responds directly to the shocks. However, in order for there to be a nontrivial inference problem for agents, we need some components of the shocks to be unobservable. Note that we’ve assumed that both the slope of the Phillips curve γ_{j_t} and the interest sensitivity σ_{j_t} vary with the mode j_t . For the former, this could reflect changes in the degree of monopolistic competition (which also lead to varying markups) and/or changes in the degree of price stickiness. The interest sensitivity shift is purely a change in the preferences of the agents in the economy, although it could also result from non-homothetic preferences coupled with shifts in output (in which case there would be no shift in the preferences themselves, but the intertemporal elasticity would vary with the level of output). Unlike our illustration above, there are no switches in the steady state levels of the variables of interest here, as we consider the usual approximations around a zero inflation rate and an efficient level of output.

3.2 Optimal policy: NL, AOP, and BOP

Here we examine value functions and optimal policies for this simple New Keynesian model under no learning (NL), adaptive optimal policy (AOP), and Bayesian optimal policy (BOP). We use the following loss function,

$$L_t = \pi_t^2 + \lambda_j y_t^2 + \mu i_t^2, \quad (3.4)$$

We set the following parameters, mostly following Woodford's [21] calibration as follows: $\gamma_1 = 0.024$, $\gamma_2 = 0.075$, $\sigma_1 = 1/.157 = 6.37$, $\sigma_2 = 1$, $c_\pi = c_y = c_g = 0.5$, and $\rho = 0.5$. We set the loss function parameters as: $\delta = 0.99$, $\lambda_j = 2\gamma_j$, and $\mu = 0.236$. Most of the structural parameters are taken from Woodford [21], while the two modes represent reasonable alternatives. Mode 1 is Woodford's benchmark case, while mode 2 has a substantially smaller interest rate sensitivity (one consistent with logarithmic preferences) and a larger response γ of inflation to output. We set the transition matrix to


$$P = \begin{bmatrix} 0.98 & 0.02 \\ 0.02 & 0.98 \end{bmatrix}.$$

We have two forward-looking variables ($x_t \equiv (\pi_t, y_t)'$) and consequently two Lagrange multipliers ($\Xi_{t-1} \equiv (\Xi_{\pi,t-1}, \Xi_{y,t-1})'$). We have one predetermined variable ($X_t \equiv g_t$) and the estimated mode probabilities ($p_{t|t} \equiv (p_{1t|t}, p_{2t|t})'$) (of which we only need keep track of one, $p_{1t|t}$). Thus, the value and policy functions, $V(s_t)$ and $i(s_t)$, are all four dimensional ($s_t = (g_t, \Xi'_{t-1}, p_{1t|t})'$). Thus we are forced for computational reasons to restrict attention to relatively sparse grids with few points. The following plots show two-dimensional slices of the value and policy functions, focusing on the dependence on g_t and $p_{1t|t}$ (which we for simplicity denote by p_{1t} in the figures). In particular, all of the plots are for $\Xi_{t-1} = (0, 0)'$.

Figure 3.1 shows losses under NL and BOP as functions of p_{1t} and g_t . Figure 3.2 shows the difference between losses under NL, AOP, and BOP. Figures 3.3 and 3.4 show the corresponding policy functions and their differences.

In Svensson and Williams [17] we show that learning implies a mean-preserving spread of the random variable $p_{t+1|t+1}$ (which is under learning a random variable from the vantage point of period t). Hence, concavity of the value function under NL in p_{1t} implies that learning is beneficial, since then a mean-preserving spread reduces the expected future loss. However, we see in figure 3.1 that the value function is actually slightly convex in p_{1t} , so learning is not beneficial here. In contrast, for a backward-looking example in Svensson and Williams [17], the value function is concave and learning is beneficial.


Figure 3.1: Losses from no learning (NL) and Bayesian optimal policy (BOP)



Consequently, we see in figure 3.2 that AOP gives higher losses than NL. Furthermore, somewhat surprisingly, we see that BOP gives higher losses than AOP (although the difference is very small). This is all counter to an example with a backward-looking model in Svensson and Williams [17].

Why is this different in a model with forward-looking variables? It may at least partially be a remnant of our assumption of symmetric beliefs and information between the private sector and the policymaker. With backward-looking models, we have generally found that learning is beneficial. Moreover, with backward-looking models, the BOP is always weakly better than the AOP, as acknowledging the endogeneity of information in the BOP case need not mean that policy must change. (That is, the AOP policy is always feasible in the BOP problem.) However, with forward-looking models, neither of these conclusions holds. Under our assumption of symmetric information and beliefs between the private sector and the policymaker, both the private sector and the policymaker learns. The difference then comes from the way that private sector beliefs also respond to learning and to the experimentation motive. Having more reactive private sector beliefs may add volatility and make it more difficult for the policymaker to stabilize the economy. Acknowledging the endogeneity of information in the BOP case then need not be beneficial either,

Figure 3.2: Differences in losses from no learning (NL), adaptive optimal policy (AOP) and Bayesian optimal policy (BOP)




as it may induce further volatility in agents' beliefs. (Note that, in the forward-looking case, we solve saddlepoint problems, and in going from AOP to BOP we are expanding the feasible set for both the minimizing and maximizing choices.)

4 Learning in an estimated empirical New Keynesian model

In the previous section we focused on a simple small model in order to consider the impacts of learning and experimentation. As computing BOP is computationally intensive, there are limits to the degree of empirical realism of the models we can address in that framework. In this section we focus on a more empirically plausible model, a version of the model of Lindé [11] that we estimated in Svensson and Williams [16]. This model includes richer dynamics for inflation and the output gap, which both have backward and forward-looking components. However, these additional dynamics increase the dimension of the state space, which implies that it is not very feasible to consider the BOP. Thus we focus here on the impact of learning on policy and compare NL and AOP. In Svensson and Williams [16] we computed the optimal policy under no-learning, and here

Figure 3.3: Optimal policies under no learning (NL) and Bayesian optimal policy (BOP)



we see how inference on the mode affects the dynamics of output, inflation, and interest rates.

4.1 The model

The structural model is a mode-dependent simplification of the model of the US economy of Lindé [11] and is given by

$$\begin{aligned} \pi_t &= \omega_{fj} E_t \pi_{t+1} + (1 - \omega_{fj}) \pi_{t-1} + \gamma_j y_t + c_{\pi j} \varepsilon_{\pi t}, \\ y_t &= \beta_{fj} E_t y_{t+1} + (1 - \beta_{fj}) [\beta_{yj} y_{t-1} + (1 - \beta_{yj}) y_{t-2}] - \beta_{rj} (i_t - E_t \pi_{t+1}) + c_{yj} \varepsilon_{y t}. \end{aligned} \quad (4.1)$$

Here $j \in \{1, 2\}$ indexes the mode, and the shocks $\varepsilon_{\pi t}$, $\varepsilon_{y t}$, and ε_{it} are independent standard normal random variables. In particular, we consider a two-mode MJLQ model where one mode has forward- and backward-looking elements, while the other is backward-looking only. Thus we specify that mode 1 is unrestricted, while in mode 2 we restrict $\omega_f = \beta_f = 0$, so that the mode is backward-looking. For estimation, we also impose a particular instrument rule for i_t , but as we focus on optimal policy we do not include that here.

In Svensson and Williams [16] we estimate the model on US data using Bayesian methods, with the maximum posterior estimates given in table 4.1, with the unconditional expectation of

Figure 3.4: Differences in policies under no learning (NL) and Bayesian optimal policy (BOP)




Table 4.1: Estimates of the constant-coefficient and a restricted two-mode Lindé model.


Parameter	Mean	Mode 1	Mode 2
ω_f	0.0938	0.3272	0
γ	0.0474	0.0580	0.0432
β_f	0.1375	0.4801	0
β_r	0.0304	0.0114	0.0380
β_y	1.3331	1.5308	1.2538
c_π	0.8966	1.0621	0.8301
c_y	0.5572	0.5080	0.5769

the coefficients for comparison. Here we see that apart from the forward-looking terms (which of course are restricted) the variation in the other parameters across the modes is relatively minor. There are some differences in the estimated policy functions (not reported here), but relatively little change across modes in the other structural coefficients. The estimated transition matrix P and its implied stationary distribution \bar{p} are given by

$$P = \begin{bmatrix} 0.9579 & 0.0421 \\ 0.0169 & 0.9831 \end{bmatrix}, \quad \bar{p} = \begin{bmatrix} 0.2869 \\ 0.7131 \end{bmatrix}.$$

Thus mode 2 is the most persistent and has the largest mass in the invariant distribution. This is consistent with our estimation of the modes as shown in figure 4.1. Again, the plots show both the smoothed and filtered estimates. Mode 2, the backward-looking model mode, was experienced the most throughout much of the sample, holding for 1961–1968 and then with near certainty continually since 1985. The forward-looking model held in periods of rapid changes in inflation, holding for both the run-ups in inflation in the early and late 1970s and the disinflationary period

Figure 4.1: Estimated probabilities of being the different modes. Solid lines: smoothed (full-sample) inference. Dashed lines: filtered (one-sided) inference.



of the early 1980s. During periods of relative tranquility, such as the Greenspan era, the backward-looking model fits the data the best.

4.2 Optimal policy: NL and AOP


Using the methods described above, we solve for the optimal policy functions

$$i_t = F_i(p_t | t) \tilde{X}_t,$$

where now $\tilde{X}_t \equiv (\pi_{t-1}, y_{t-1}, y_{t-2}, i_{t-1}, \Xi_{\pi, t-1}, \Xi_{y, t-1})'$. In Svensson and Williams [16] we focused on the observable and no-learning cases, and assumed that the shocks $\varepsilon_{\pi t}$ and $\varepsilon_{y t}$ were observable. Thus we set $C_2 \equiv 0$ and treated the shocks as additional predetermined variables. However, to focus on the role of learning, we now assume that those shocks are unobservable. If they were observable, then agents would be able to infer the mode from their observations of the forward-looking variables. We use the following loss function:

$$L_t = \pi_t^2 + \lambda y_t^2 + \nu (i_t - i_{t-1})^2, \quad (4.2)$$

Figure 4.2: Unconditional impulse responses to shocks under the optimal policy for the two-mode version of the Lindé model. Solid lines: median responses under AOP. Dashed lines: median responses under NL. Dot-dashed lines: constant-coefficient responses.



which is a common central-bank loss function. We set the weights to $\lambda = 1$ and $\nu = 0.2$, and fix the discount factor in the intertemporal loss function to $\delta = 1$.

For ease of interpretation, we plot the distribution of the impulse responses of inflation, the output gap, and the instrument rate to the two structural shocks in figure 4.2. We consider 10,000 simulations of 50 periods, and plot the median responses for the optimal policy under NL and AOP, and the corresponding optimal responses for the constant-coefficient model.¹³

Compared to the constant-coefficient case, the mean impulse responses are consistent with larger effects of the shocks that are also longer lasting. In terms of the optimal policy responses, the AOP and NL cases are quite similar, and in both cases the peak response to shocks is nearly the same as in the constant-coefficient case, but it comes with a delay. Again compared to the constant-coefficient case, the responses of inflation and the output gap are larger and more sustained when there is model uncertainty.

¹³ The shocks are $\varepsilon_{\pi 0} = 1$ and $\varepsilon_{y 0} = 1$, respectively, so the shocks to the inflation and output-gap equations in period 0 are mode-dependent and equal to $c_{\pi j}$ and $c_{y j}$ ($j = 1, 2, 3$), respectively. The distribution of modes in period 0 (and thereby all periods) is again the stationary distribution.

Figure 4.3: Distribution (across samples) of various statistics under the optimal policy for the two-mode version of the Lindé model. Solid lines: AOP. Dashed lines: NL.





Table 4.2: Average of different statistics from 1000 simulations of 1000 periods each of our estimated model under the no-learning (NL) and adaptive (AOP) optimal policies.

Policy	$E \pi_t$	Std π_t	$E y_t$	Std y_t	$E i_t$	Std i_t	$E L_t$
NL	-0.1165	5.2057	0.1303	5.6003	0.0073	10.0239	88.4867
AOP	-0.0300	3.1696	0.0299	2.7698	0.0011	9.9989	38.8710

However, here we see that learning can be beneficial, as the optimal policy under AOP dampens the responses to shocks, particularly for shocks to inflation. As the optimal policy responses are nearly identical, this seems to be largely due to more accurate forecasts by the public, which lead to more rapid stabilization.

While these impulse responses are revealing, they do not capture the full benefits from learning, as by definition they simply provide the responses to a single shock. To gain a better understanding of the role of learning, we now simulate our model under the NL and AOP policies to compare the realized economic performance. Table 4.2 summarizes various statistics resulting from 1000 simulations of 1000 periods each. Thus for example, the entry there under “ $E\pi_t$ ” is the average

Figure 4.4: Simulated time series under the optimal policy for the two-mode version of the Lindé model. *Top three panels:* Solid lines: AOP. Dashed lines: NL. *Bottom panel:* Solid line: probability of mode 1. Dotted line: true mode. Dashed line: unconditional probability of mode 1.



across the 1000 simulations of the sample average (over the 1000 periods) of inflation, while “Std π_t ” is the average across simulations of the standard deviation (in each time series) of inflation. In particular, we see from the entry under “ EL_t ” that the average period loss is less than half under AOP compared to NL. In addition to these averages, figure 4.3 plots the distribution across samples of the key components of the loss function. There we plot a kernel smoothed estimate of the distribution from the 1000 simulations. We see that the distribution of sample losses is much more favorable under AOP than under NL.

In figure 4.4 we show one representative simulation to illustrate the differences. The more effective stabilization of inflation and the output gap under AOP for very similar instrument-rate settings as under NL is apparent.

5 Conclusions

In this paper, we have presented a relatively general framework for analyzing model uncertainty and the interactions between learning and optimization. While this is a classic issue, very little to date has been done for systems with forward-looking variables, which are essential elements of modern models for policy analysis. Our specification is general enough to cover many practical cases of interest, but yet remains relatively tractable in implementation. This is definitely true for cases when decision makers do not learn from the data they observe (our case of no learning, NL) or when they do learn but do not account for learning in optimization (our case of adaptive optimal policy, AOP). In both of these cases, we have developed efficient algorithms for solving for the optimal policy, which can handle relatively large models with multiple modes and many state variables. However, in the case of the Bayesian optimal policy (BOP), where the experimentation motive is taken into account, we must solve more complex numerical dynamic programming problems. Thus to fully examine optimal experimentation we are haunted by the curse of dimensionality, forcing us to study relatively small and simple models.

Thus, an issue of much practical importance is the size of the experimentation component of policy, and the losses entailed by abstracting from it. While our results in this paper are far from comprehensive, they suggest that in practical settings the experimentation motive may not be a concern. The above and similar examples that we have considered indicate that the benefits of learning (moving from NL to AOP) may be substantial, whereas the benefits from experimentation (moving from AOP to BOP) are modest or even insignificant. If this preliminary finding stands up to scrutiny, experimentation in economic policy in general and monetary policy in particular may not be very beneficial, in which case there is little need to face the difficult ethical and other issues involved in conscious experimentation in economic policy. Furthermore, the AOP is much easier to compute and implement than the BOP. To have this truly be a robust implication, more simulations and cases need to be examined.

References

- [1] Beck, Günter W., and Volker Wieland (2002), “Learning and Control in a Changing Economic Environment,” *Journal of Economic Dynamics and Control* 25, 1359–1377.

- [2] Benigno, Piarpalo, and Michael Woodford (2007), “Linear-Quadratic Approximation of Optimal Policy Problems,” working paper, Columbia University.
- [3] Blake, Andrew P., and Fabrizio Zampolli (2006), “Time Consistent Policy in Markov Switching Models,” Working Paper No. 298, Bank of England.
- [4] Cogley, Timothy, Riccardo Colacito, and Thomas J. Sargent (2007) “The Benefits from U.S. Monetary Policy Experimentation in the Days of Samuelson and Solow and Lucas,” *Journal of Money Credit and Banking* 39, s67–99.
- [5] Costa, Oswaldo L.V., Marecelo D. Fragoso, and Ricardo P. Marques (2005), *Discrete-Time Markov Jump Linear Systems*, Springer, London.
- [6] do Val, João B.R., and Tamer Başar (1999), “Receding Horizon Control of Jump Linear Systems and a Macroeconomic Policy Problem,” *Journal of Economic Dynamics and Control* 23, 1099–1131.
- [7] Ellison, Martin (2006), “The Learning Cost of Interest Rate Reversals,” *Journal of Monetary Economics* 53, 1895–1907.
- [8] Ellison, Martin, and Natacha Valla (2001), “Learning, Uncertainty and Central Bank Activism in an Economy with Strategic Interactions,” *Journal of Monetary Economics* 48, 153–171.
- [9] Evans, George, and Seppo Honkapohja (2001) *Learning and Expectations in Macroeconomics*, Princeton University Press, Princeton.
- [10] Kiefer, Nicholas M. (1989), “A Value Function Arising in the Economics of Information,” *Journal of Economic Dynamics and Control* 13, 201–223.
- [11] Lindé, Jesper (2005), “Estimating New-Keynesian Phillips Curves: A Full Information Maximum Likelihood Approach,” *Journal of Monetary Economics* 52, 1135–1149.
- [12] Marcet, Albert, and Ramon Marimon (1998), “Recursive Contracts,” working paper, www.econ.upf.edu.
- [13] Miranda, Mario J., and Paul L. Fackler (2002), *Applied Computational Economics and Finance*, MIT Press.

- [14] Rudebusch, Glenn D., and Lars E.O. Svensson (1999), “Policy Rules for Inflation Targeting,” in John B. Taylor (ed.), *Monetary Policy Rules*, University of Chicago Press.
- [15] Svensson, Lars E.O. (2007), “Optimization under Commitment and Discretion, the Recursive Saddlepoint Method, and Targeting Rules and Instrument Rules: Lecture Notes,” www.princeton.edu/svensson.
- [16] Svensson, Lars E.O., and Noah Williams (2007a), “Monetary Policy with Model Uncertainty: Distribution Forecast Targeting,” working paper, www.princeton.edu/svensson/ and www.princeton.edu/~noahw/.
- [17] Svensson, Lars E.O., and Noah Williams (2007b), “Bayesian and Adaptive Optimal Policy Under Model Uncertainty,” working paper, www.princeton.edu/svensson/ and www.princeton.edu/~noahw/.
- [18] Tesfaselassie, Mewael F., Eric Schaling, and Sylvester C. W. Eijffinger (2006), “Learning about the Term Structure and Optimal Rules for Inflation Targeting,” CEPR Discussion paper No. 5896.
- [19] Wieland, Volker (2000), “Learning by Doing and the Value of Optimal Experimentation,” *Journal of Economic Dynamics and Control* 24, 501–534.
- [20] Wieland, Volker (2006), “Monetary Policy and Uncertainty about the Natural Unemployment Rate: Brainard-Style Conservatism versus Experimental Activism,” *Advances in Macroeconomics* 6(1).
- [21] Woodford, Michael (2003), *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton University Press.
- [22] Zampolli, Fabrizio (2006), “Optimal Monetary Policy in a Regime-Switching Economy: The Response to Abrupt Shifts in Exchange-Rate Dynamics,” Working Paper No. 297, Bank of England.